Strategic games

Krzysztof R. Apt

May 18, 2011

Introduction

Mathematical game theory, as launched by Von Neumann and Morgenstern in their seminal book von Neumann and Morgenstern [1944], followed by Nash' contributions Nash [1950,1951], has become a standard tool in Economics for the study and description of various economic processes, including competition, cooperation, collusion, strategic behaviour and bargaining. Since then it has also been successfuly used in Biology, Political Sciences, Psychology and Sociology. With the advent of the Internet game theory became increasingly relevant in Computer Science.

One of the main areas in game theory are ***strategic games***, (sometimes also called ***non-cooperative games***), which form a simple model of interaction between profit maximizing players. In strategic games each player has a payoff function that he aims to maximize and the value of this function depends on the decisions taken *simultaneously* by all players. Such a simple description is still amenable to various interpretations, depending on the assumptions about the existence of *private information.* The purpose of these lecture notes is to provide a simple introduction to the most common concepts used in strategic games and most common types of such games.

Many books provide introductions to various areas of game theory, including strategic games. Most of them are written from the perspective of applications to Economics. In the nineties the leading textbooks were Myerson [1991], Binmore [1991], Fudenberg and Tirole [1991] and Osborne and Rubinstein [1994].

Moving to the next decade, Osborne [2005] is an excellent, broad in its scope, undergraduate level textbook, while Peters [2008] is probably the best book on the market for the graduate level. Undeservedly less known is the short and lucid Tijs [2003]. An elementary, short introduction, focusing on the concepts, is Shoham and Leyton-Brown [2008]. In turn, Ritzberger [2002] is a comprehensive book on strategic games that also extensively discusses ***extensive games***, i.e., games in which the players choose actions in turn. Finally, Binmore [2007] is thoroughly revised version of Binmore [1991].

Several textbooks on microeconomics include introductory chapters on game theory, including strategic games. Two good examples are Mas-Collel, Whinston and Green [1995] and Jehle and Reny [2000]. Finally, Nisan et al. [2007] is a recent collection of surveys and introductions to the computational aspects of game theory, with a number of articles concerned with strategic games and mechanism design.

Chapter 1

Nash Equilibrium

Assume a set $\{1, \ldots, n\}$ of players, where $n > 1$. A ***strategic game*** (or ***non-cooperative game***) for n players, written as $(S_1, \ldots, S_n, p_1, \ldots, p_n)$, consists of

- a non-empty (possibly infinite) set S_i of ***strategies***,
- a ***payoff function*** $p_i : S_1 \times \ldots \times S_n \to \mathbb{R}$,

for each player i.

We study strategic games under the following basic assumptions:

- players choose their strategies *simultaneously*; subsequently each player receives a payoff from the resulting joint strategy,
- each player is ***rational***, which means that his objective is to maximize his payoff,
- players have ***common knowledge*** of the game and of each others' rationality.[1]

Here are three classic examples of strategic two-player games to which we shall return in a moment. We represent such games in the form of a bimatrix, the entries of which are the corresponding payoffs to the row and column players. So for instance in the Prisoner's Dilemma game, when the row player chooses C (cooperate) and the column player chooses D (defect),

[1] Intuitively, common knowledge of some fact means that everybody knows it, everybody knows that everybody knows it, etc. This notion can be formalized using epistemic logic.

then the payoff for the row player is 0 and the payoff for the column player is 3.

Prisoner's Dilemma

	C	D
C	$2, 2$	$0, 3$
D	$3, 0$	$1, 1$

Battle of the Sexes

	F	B
F	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

Matching Pennies

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

We introduce now some basic notions that will allow us to discuss and analyze strategic games in a meaningful way. Fix a strategic game

$$(S_1, \ldots, S_n, p_1, \ldots, p_n).$$

We denote $S_1 \times \ldots \times S_n$ by S, call each element $s \in S$ a ***joint strategy***, or a ***strategy profile***, denote the ith element of s by s_i, and abbreviate the sequence $(s_j)_{j \neq i}$ to s_{-i}. Occasionally we write (s_i, s_{-i}) instead of s. Finally, we abbreviate $\times_{j \neq i} S_j$ to S_{-i} and use the '$_{-i}$' notation for other sequences and Cartesian products.

We call a strategy s_i of player i a ***best response*** to a joint strategy s_{-i} of his opponents if

$$\forall s'_i \in S_i \ p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

Next, we call a joint strategy s a ***Nash equilibrium*** if each s_i is a best response to s_{-i}, that is, if

$$\forall i \in \{1, \ldots, n\} \ \forall s'_i \in S_i \ p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

So a joint strategy is a Nash equilibrium if no player can achieve a higher payoff by *unilaterally* switching to another strategy. Intuitively, a Nash equilibrium is a situation in which each player is a posteriori satisfied with his choice.

Let us return now the three above introduced games.

Re: Prisoner's Dilemma

The Prisoner's Dilemma game has a unique Nash equilibrium, namely (D, D). One of the peculiarities of this game is that in its unique Nash equilibrium each player is worse off than in the outcome (C, C). We shall return to this game once we have more tools to study its characteristics.

To clarify the importance of this game we now provide a couple of simple interpretations of it. The first one, due to Aumann, is the following.

> Each player decides whether he will receive 1000 dollars or the other will receive 2000 dollars. The decisions are simultaneous and independent.

So the entries in the bimatrix of the Prisoner's Dilemma game refer to the thousands of dollars each player will receive. For example, if the row player asks to give 2000 dollars to the other player, and the column player asks for 1000 dollar for himself, the row player gets nothing while column player gets 3000 dollars. This contingency corresponds to the 0,3 entry in the bimatrix.

The original interpretation of this game that explains its name refers to the following story.

> Two suspects are taken into custody and separated. The district attorney is certain that they are guilty of a specific crime, but he does not have adequate evidence to convict them at a trial. He points out to each prisoner that each has two alternatives: to confess to the crime the police are sure they have done (C), or not to confess (N).
>
> If they both do not confess, then the district attorney states he will book them on some very minor trumped-up charge such as petty larceny or illegal possession of weapon, and they will both receive minor punishment; if they both confess they will be prosecuted, but he will recommend less than the most severe sentence; but if one confesses and the other does not, then the confessor

will receive lenient treatment for turning state's evidence whereas the latter will get "the book" slapped at him.

This is represented by the following bimatrix, in which each negative entry, for example -1, corresponds to the 1 year prison sentence ('the lenient treatment' referred to above):

	C	N
C	$-5, -5$	$-1, -8$
N	$-8, -1$	$-2, -2$

The negative numbers are used here to be compatible with the idea that each player is interested in maximizing his payoff, so, in this case, of receiving a lighter sentence. So for example, if the row suspect decides to confess, while the column suspect decides not to confess, the row suspect will get 1 year prison sentence (the 'lenient treatment'), the other one will get 8 years of prison (' "the book" slapped at him').

Many other natural situations can be viewed as a Prisoner's Dilemma game. This allows us to explain the underlying, undesidered phenomena.

Consider for example the arms race. For each of two warring, equally strong countries, it is beneficial not to arm instead of to arm. Yet both countries end up arming themselves. As another example consider a couple seeking a divorce. Each partner can choose an inexpensive (bad) or an expensive (good) layer. In the end both partners end up choosing expensive lawyers. Next, suppose that two companies produce a similar product and may choose between low and high advertisement costs. Both end up heavily advertising.

Re: Matching Pennies game

Next, consider the Matching Pennies game. This game formalizes a game that used to be played by children. Each of two children has a coin and simultaneously shows heads (H) or tails (T). If the coins match then the first child wins, otherwise the second child wins. This game has no Nash equilibrium. This corresponds to the intuition that for no outcome both players are satisfied. Indeed, in each outcome the losing player regrets his choice. Moreover, the social welfare of each outcome is 0. Such games are called ***zero sum games*** and we shall return to them later. Also, we shall return to this game once we have introduced mixed strategies.

Re: Battle of the Sexes game

Finally, consider the Battle of the Sexes game. The interpretation of this game is as follows. A couple has to decide whether to go out for a football match (F) or a ballet (B). The man, the row player prefers a football match over the ballet, while the woman, the column player, the other way round. Moreover, each of them prefers to go out together than to end up going out separately. This game has two Nash equilibria, namely (F, F) and (B, B). Clearly, there is a problem how the couple should choose between these two satisfactory outcomes. Games of this type are called ***coordination games***.

Obviously, all three games are very simplistic. They deal with two players and each player has to his disposal just two strategies. In what follows we shall introduce many interesting examples of strategic games. Some of them will deal with many players and some games will have several, sometimes an infinite number of strategies.

To close this chapter we consider two examples of more interesting games, one for two players and another one for an arbitrary number of players.

Example 1 (Traveler's dilemma)

Suppose that two travellers have identical luggage, for which they both paid the same price. Their luggage is damaged (in an identical way) by an airline. The airline offers to recompense them for their luggage. They may ask for any dollar amount between \$2 and \$100. There is only one catch. If they ask for the same amount, then that is what they will both receive. However, if they ask for different amounts —say one asks for $\$m$ and the other for $\$m'$, with $m < m'$— then whoever asks for $\$m$ (the lower amount) will get $\$(m + 2)$, while the other traveller will get $\$(m - 2)$. The question is: what amount of money should each traveller ask for?

We can formalize this problem as a two-player strategic game, with the set $\{2, \ldots, 100\}$ of natural numbers as possible strategies. The following payoff function[2] formalizes the conditions of the problem:

$$p_i(s) := \begin{cases} s_i & \text{if } s_i = s_{-i} \\ s_i + 2 & \text{if } s_i < s_{-i} \\ s_{-i} - 2 & \text{otherwise} \end{cases}$$

It is easy to check that $(2, 2)$ is a Nash equilibrium. To check for other Nash equilibria consider any other combination of strategies (s_i, s_{-i}) and

[2]We denote in two-player games the opponent of player i by $-i$, instead of $3 - i$.

suppose that player i submitted a larger or equal amount, i.e., $s_i \geq s_{-i}$. Then player's i payoff is s_{-i} if $s_i = s_{-i}$ or $s_{-i} - 2$ if $s_i > s_{-i}$.

In the first case he will get a strictly higher payoff, namely $s_{-i} + 1$, if he submits instead the amount $s_{-i} - 1$. (Note that $s_i = s_{-i}$ and $(s_i, s_{-i}) \neq (2, 2)$ implies that $s_{-i} - 1 \in \{2, \ldots, 100\}$.) In turn, in the second case he will get a strictly higher payoff, namely s_{-i}, if he submits instead the amount s_{-i}.

So in each joint strategy $(s_i, s_{-i}) \neq (2, 2)$ at least one player has a strictly better alternative, i.e., his strategy is not a best response. This means that $(2, 2)$ is a unique Nash equilibrium. This is a paradoxical conclusion, if we recall that informally a Nash equilibrium is a state in which both players are satisfied with their choice. □

Example 2 Consider the following ***beauty contest game***. In this game there are $n > 2$ players, each with the set of strategies equal $\{1, \ldots, 100\}$, Each player submits a number and the payoff to each player is obtained by splitting 1 equally between the players whose submitted number is closest to $\frac{2}{3}$ of the average. For example, if the submissions are $29, 32, 29$, then the payoffs are respectively $\frac{1}{2}, 0, \frac{1}{2}$.

Finding Nash equilibria of this game is not completely straightforward. At this stage we only observe that the joint strategy $(1, \ldots, 1)$ is clearly a Nash equilibrium. We shall answer the question of whether there are more Nash equilibria once we introduce some tools to analyze strategic games. □

Exercise 1 Find all Nash equilibria in the following games:

Stag hunt

	S	R
S	$2, 2$	$0, 1$
R	$1, 0$	$1, 1$

Coordination

	L	R
T	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Pareto Coordination

	L	R
T	$2, 2$	$0, 0$
B	$0, 0$	$1, 1$

Hawk-dove

	H	D
H	$0,0$	$3,1$
D	$1,3$	$2,2$

$\square$

Exercise 2 Consider the following ***inspection game***.

There are two players: a worker and the boss. The worker can either Shirk or put an Effort, while the boss can either Inspect or Not. Finding a shirker has a benefit b while the inspection costs c, where $b > c > 0$. So if the boss carries out an inspection his benefit is $b - c > 0$ if the worker shirks and $-c < 0$ otherwise.

The worker receives 0 if he shirks and is inspected, and g if he shirks and is not found. Finally, the worker receives w, where $g > w > 0$ if he puts in the effort.

This leads to the following bimatrix:

	I	N
S	$0, b-c$	$g, 0$
E	$w, -c$	$w, 0$

Analyze the best responses in this game. What can we conclude from it about the Nash equilibria of this game?

$\square$

Chapter 2

Nash Equilibria and Pareto Efficient Outcomes

To discuss strategic games in a meaningful way we need to introduce further, natural, concepts. Fix a strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.

We call a joint strategy s a ***Pareto efficient outcome*** if for no joint strategy s'

$$\forall i \in \{1, \ldots, n\}\ p_i(s') \geq p_i(s) \text{ and } \exists i \in \{1, \ldots, n\}\ p_i(s') > p_i(s).$$

That is, a joint strategy is a Pareto efficient outcome if no joint strategy is both a weakly better outcome for all players and a strictly better outcome for some player.

Further, given a joint strategy s we call the sum $\sum_{j=1}^{n} p_j(s)$ the ***social welfare*** of s. Next, we call a joint strategy s a ***social optimum*** if the social welfare of s is maximal.

Clearly, if s is a social optimum, then s is Pareto efficient. The converse obviously does not hold. Indeed, in the Prisoner's Dilemma game the joint strategis (C, D) and (D, C) are both Pareto efficient, but their social welfare is not maximal. Note that (D, D) is the only outcome that is not Pareto efficient. The social optimum is reached in the strategy profile (C, C). In contrast, the social welfare is smallest in the Nash equlibrium (D, D).

This discrepancy between Nash equilibria and Pareto efficient outcomes is absent in the Battle of Sexes game. Indeed, here both concepts coincide.

The tension between Nash equlibria and Pareto efficient outcomes present in the Prisoner's Dilemma game occurs in several other natural games. It forms one of the fundamental topics in the theory of strategic games. In this chapter we shall illustrate this phenomenon by a number of examples.

Example 3 (Prisoner's Dilemma for n players)
First, the Prisoner's Dilemma game can be easily generalized to n players as follows. It is convenient to assume that each player has two strategies, 1, representing cooperation, (formerly C) and 0, representing defection, (formerly D). Then, given a joint strategy s_{-i} of the opponents of player i, $\sum_{j \neq i} s_j$ denotes the number of 1 strategies in s_{-i}. Denote by $\mathbf{1}$ the joint strategy in which each strategy equals 1 and similarly with $\mathbf{0}$.

We put

$$p_i(s) := \begin{cases} 2\sum_{j \neq i} s_j + 1 & \text{if } s_i = 0 \\ 2\sum_{j \neq i} s_j & \text{if } s_i = 1 \end{cases}$$

Note that for $n = 2$ we get the original Prisoner's Dilemma game.

It is easy to check that the strategy profile $\mathbf{0}$ is the unique Nash equilibrium in this game. Indeed, in each other strategy profile a player who chose 1 (cooperate) gets a higher payoff when he switches to 0 (defect).

Finally, note that the social welfare in $\mathbf{1}$ is $2n(n-1)$, which is strictly more than n, the social welfare in $\mathbf{0}$. We now show that $2n(n-1)$ is the social optimum. To this end it suffices to note that if a single player switches from 0 to 1, then his payoff decreases by 1 but the payoff of each other player increases by 2, and hence the social welfare increases. $\square$

The next example deals with the depletion of ***common resources***, which in economics are goods that are not *excludable* (people cannot be prevented from using them) but are *rival* (one person's use of them diminishes another person's enjoyment of it). Examples are congested toll-free roads, fish in the ocean, or the environment. The overuse of such common resources leads to their destruction. This phenomenon is called the ***tragedy of the commons***.

One way to model it is as a Prisoner's dilemma game for n players. But such a modeling is too crude as it does not reflect the essential characteristics of the problem. We provide two more adequate modeling of it, one for the case of a binary decision (for instance, whether to use a congested road or not), and another one for the case when one decides about the intensity of using the resource (for instance on what fraction of a lake should one fish).

Example 4 (Tragedy of the commons I)

Assume $n > 1$ players, each having to its disposal two strategies, 1 and 0 reflecting, respectively, that the player decides to use the common resource or not. If he does not use the resource, he gets a fixed payoff. Further, the users of the resource get the same payoff. Finally, the more users of the common resource the smaller payoff for each of them gets, and when the number of users exceeds a certain threshold it is better for the other players not to use the resource.

The following payoff function realizes these assumptions:

$$p_i(s) := \begin{cases} 0.1 & \text{if } s_i = 0 \\ F(m)/m & \text{otherwise} \end{cases}$$

where $m = \sum_{j=1}^{n} s_j$ and

$$F(m) := 1.1m - 0.1m^2.$$

Indeed, the function $F(m)/m$ is strictly decreasing. Moreover, $F(9)/9 = 0.2$, $F(10)/10 = 0.1$ and $F(11)/11 = 0$. So when there are already ten or more users of the resource it is indeed better for other players not to use the resource.

To find a Nash equilibrium of this game, note that given a strategy profile s with $m = \sum_{j=1}^{n} s_j$ player i profits from switching from s_i to another strategy in precisely two circumstances:

- $s_i = 0$ and $F(m+1)/(m+1) > 0.1$,
- $s_i = 1$ and $F(m)/m < 0.1$.

In the first case we have $m + 1 < 10$ and in the second case $m > 10$.

Hence when $n < 10$ the only Nash equilibrium is when all players use the common resource and when $n \geq 10$ then s is a Nash equilibrium when either 9 or 10 players use the common resource.

Assume now that $n \geq 10$. Then in a Nash equilibrium s the players who use the resource receive the payoff 0.2 (when $m = 9$) or 0.1 (when $m = 10$). So the maximum social welfare that can be achieved in a Nash equilibrium is $0.1(n - 9) + 1.8 = 0.1n + 0.9$.

To find a strategy profile in which social optimum is reached with the largest social welfare we need to find m for which the function $0.1(n - m) + F(m)$ reaches the maximum. Now, $0.1(n - m) + F(m) = 0.1n + m - 0.1m^2$

and by elementary calculus we find that $m = 5$ for which $0.1(n-m)+F(m) = 0.1n+2.5$. So the social optimum is achieved when 5 players use the common resource. $\square$

Example 5 (Tragedy of the commons II)
Assume $n > 1$ players, each having to its disposal an infinite set of strategies that consists of the real interval $[0, 1]$. View player's strategy as its chosen fraction of the common resource. Then the following payoff function reflects the fact that player's enjoyment of the common resource depends positively from his chosen fraction of the resource and negatively from the total fraction of the common resource used by all players:

$$p_i(s) := \begin{cases} s_i(1 - \sum_{j=1}^{n} s_j) & \text{if } \sum_{j=1}^{n} s_j \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

The second alternative reflects the phenomenon that if the total fraction of the common resource by all players exceeds a feasible level, here 1, then player's enjoyment of the resource becomes zero. We can write the payoff function in a more compact way as

$$p_i(s) := \max(0, s_i(1 - \sum_{j=1}^{n} s_j)).$$

To find a Nash equilibrium of this game, fix $i \in \{1, \ldots, n\}$ and s_{-i} and denote $\sum_{j \neq i} s_j$ by t. Then $p_i(s_i, s_{-i}) = \max(0, s_i(1 - t - s_i))$.

By elementary calculus player's i payoff becomes maximal when $s_i = \frac{1-t}{2}$. This implies that when for all $i \in \{1, \ldots, n\}$ we have

$$s_i = \frac{1 - \sum_{j \neq i} s_j}{2},$$

then s is a Nash equilibrium. This system of n linear equations has a unique solution $s_i = \frac{1}{n+1}$ for $i \in \{1, \ldots, n\}$. In this strategy profile each player's payoff is $\frac{1-n/(n+1)}{n+1} = \frac{1}{(n+1)^2}$, so its social welfare is $\frac{n}{(n+1)^2}$.

There are other Nash equilibria. Indeed, suppose that for all $i \in \{1, \ldots, n\}$ we have $\sum_{j \neq i} s_j \geq 1$, which is the case for instance when $s_i = \frac{1}{n-1}$ for $i \in \{1, \ldots, n\}$. It is straightforward to check that each such strategy profile is a Nash equilibrium in which each player's payoff is 0 and hence the social welfare is also 0. It is easy to check that no other Nash equilibria exist.

To find a strategy profile in which social optimum is reached fix a strategy profile s and let $t := \sum_{j=1}^{n} s_j$. First note that if $t > 1$, then the social welfare is 0. So assume that $t \leq 1$. Then $\sum_{j=1}^{n} p_j(s_j) = t(1-t)$. By elementary calculus this expression becomes maximal precisely when $t = \frac{1}{2}$ and then it equals $\frac{1}{4}$.

Now, for all $n > 1$ we have $\frac{n}{(n+1)^2} < \frac{1}{4}$. So the social welfare of each solution for which $\sum_{j=1}^{n} s_j = \frac{1}{2}$ is superior to the social welfare of the Nash equilibria. In particular, no such strategy profile is a Nash equilibrium.

In conclusion, the social welfare is maximal, and equals $\frac{1}{4}$, when precisely half of the common resource is used. In contrast, in the 'best' Nash equilibrium the social welfare is $\frac{n}{(n+1)^2}$ and the fraction $\frac{n}{n+1}$ of the common resource is used. So when the number of players increases, the social welfare of the best Nash equilibrium becomes arbitrarily small, while the fraction of the common resource being used becomes arbitrarily large. $\square$

The analysis carried out in the above two examples reveals that for the adopted payoff functions the common resource will be overused, to the detriment of the players (society). The same conclusion can be drawn for a much larger of class payoff functions that properly reflect the characteristics of using a common resource.

Example 6 (Cournot competition)
This example deals with a situation in which n companies independently decide their production levels of a given product. The price of the product is a linear function that depends negatively on the total output.

We model it by means of the following strategic game. We assume that for each player i:

- his strategy set is $\mathbb{R}_+$,
- his payoff function is defined by

$$p_i(s) := s_i(a - b\sum_{j=1}^{n} s_j) - cs_i$$

for some given a, b, c, where $a > c$ and $b > 0$.

Let us explain this payoff function. The price of the product is represented by the expression $a - b\sum_{j=1}^{n} s_j$, which, thanks to the assumption $b > 0$, indeed depends negatively on the total output, $\sum_{j=1}^{n} s_j$. Further, cs_i is the production cost corresponding to the production level s_i. So we assume for simplicity that the production cost functions are the same for all companies.

Further, note that if $a \leq c$, then the payoffs would be always negative or zero, since $p_i(s) = (a-c)s_i - bs_i \sum_{j=1}^{n} s_j$. This explains the assumption that $a > c$. For simplicity we do allow a possibility that the prices are negative, but see Exercise 4. The assumption $c > 0$ is obviously meaningful but not needed.

To find a Nash equilibrium of this game fix $i \in \{1,\ldots,n\}$ and s_{-i} and denote $\sum_{j \neq i} s_j$ by t. Then $p_i(s_i, s_{-i}) = s_i(a - c - bt - bs_i)$. By elementary calculus player's i payoff becomes maximal iff

$$s_i = \frac{a-c}{2b} - \frac{t}{2}.$$

This implies that s is a Nash equilibrium iff for all $i \in \{1,\ldots,n\}$

$$s_i = \frac{a-c}{2b} - \frac{\sum_{j \neq i} s_j}{2}.$$

One can check that this system of n linear equations has a unique solution, $s_i = \frac{a-c}{(n+1)b}$ for $i \in \{1,\ldots,n\}$. So this is a unique Nash equilibrium of this game.

Note that for these values of s_is the price of the product is

$$a - b\sum_{j=1}^{n} s_j = a - b\frac{n(a-c)}{(n+1)b} = \frac{a+nc}{n+1}.$$

To find the social optimum let $t := \sum_{j=1}^{n} s_j$. Then $\sum_{j=1}^{n} p_j(s) = t(a - c - bt)$. By elementary calculus this expression becomes maximal precisely when $t = \frac{a-c}{2b}$. So s is a social optimum iff $\sum_{j=1}^{n} s_j = \frac{a-c}{2b}$. The price of the product in a social optimum is $a - b\frac{a-c}{2b} = \frac{a+c}{2}$.

Now, the assumption $a > c$ implies that $\frac{a+c}{2} > \frac{a+nc}{n+1}$. So we see that the price in the social optimum is strictly higher than in the Nash equilibrium. This can be interpreted as a statement that the competition between the producers of the product drives its price down, or alternatively, that the cartel between the producers leads to higher profits for them (i.e., higher social welfare), at the cost of a higher price. So in this example reaching the social optimum is not a desirable state of affairs. The reason is that in our analysis we focussed only on the profits of the producers and omitted the customers.

As an aside also notice that when n, so the number of companies, increases, the price $\frac{a+nc}{n+1}$ in the Nash equilibrium decreases. This corresponds to the intuition that increased competition is beneficial for the customers. Note also that in the limit the price in the Nash equilibrium converges to the production cost c. $\square$

While the last two examples refer to completely different scenarios, their mathematical analysis is very similar. Their common characteristics is that in both examples the payoff functions can be written as $f(s_i, \sum_{j=1}^{n} s_j)$, where f is increasing in the first argument and decreasing in the second argument.

Exercise 3 Prove that in the game discussed in Example 5 indeed no other Nash equilibria exist apart of the mentioned ones. $\square$

Exercise 4 Modify the game from Example 6 by considering the following payoff functions:

$$p_i(s) := s_i \max(0, a - b\sum_{j=1}^{n} s_j) - cs_i.$$

Compute the Nash equilibria of this game.
Hint. Proceed as in Example 5. $\square$

Chapter 3

Strict Dominance

Let us return now to our analysis of an arbitrary strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. Let s_i, s'_i be strategies of player i. We say that s_i ***strictly dominates*** s'_i (or equivalently, that s'_i is ***strictly dominated by*** s_i) if

$$\forall s_{-i} \in S_{-i} \; p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

Further, we say that s_i is ***strictly dominant*** if it strictly dominates all other strategies of player i.

Clearly, a rational player will not choose a strictly dominated strategy. As an illustration let us return to the Prisoner's Dilemma. In this game for each player C (cooperate) is a strictly dominated strategy. So the assumption of players' rationality implies that each player will choose strategy D (defect). That is, we can predict that rational players will end up choosing the joint strategy (D, D) in spite of the fact that the Pareto efficient outcome (C, C) yields for each of them a strictly higher payoff.

The same holds in the Prisoner's Dilemma game for n players, where for all players i strategy 1 is strictly dominated by strategy 0, since for all $s_{-i} \in S_{-i}$ we have $p_i(0, s_{-i}) - p_i(1, s_{-i}) = 1$.

We assumed that each player is rational. So when searching for an outcome that is optimal for all players we can safely remove strategies that are strictly dominated by some other strategy. This can be done in a number of ways. For example, we could remove all or some strictly dominated strategies simultaneously, or start removing them in a round Robin fashion starting with, say, player 1. To discuss this matter more rigorously we introduce the notion of a restriction of a game.

Given a game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ and (possibly empty) sets of strategies $R_1, \ldots, R_n$ such that $R_i \subseteq S_i$ for $i \in \{1, \ldots, n\}$ we say that $R := (R_1, \ldots, R_n, p_1, \ldots, p_n)$ is a ***restriction*** of G. Here of course we view each p_i as a function on the subset $R_1 \times \ldots \times R_n$ of $S_1 \times \ldots \times S_n$.

In what follows, given a restriction R we denote by R_i the set of strategies of player i in R. Further, given two restrictions R and R' of G we write $R' \subseteq R$ when $\forall i \in \{1, \ldots, n\}\ R'_i \subseteq R_i$. We now introduce the following notion of reduction between the restrictions R and R' of G:

$$R \rightarrow_S R'$$

when $R \neq R'$, $R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\}\ \forall s_i \in R_i \setminus R'_i\ \exists s'_i \in R_i\ s_i \text{ is strictly dominated in } R \text{ by } s'_i.$$

That is, $R \rightarrow_S R'$ when R' results from R by removing from it some strictly dominated strategies.

We now clarify whether a one-time elimination of (some) strictly dominated strategies can affect Nash equilibria.

Lemma 1 (Strict Elimination) *Given a strategic game G consider two restrictions R and R' of G such that $R \rightarrow_S R'$. Then*

(i) if s is a Nash equilibrium of R, then it is a Nash equilibrium of R',

(ii) if G is finite and s is a Nash equilibrium of R', then it is a Nash equilibrium of R.

At the end of this chapter we shall clarify why in (ii) the restriction to finite games is necessary.

Proof.
(i) For each player the set of his strategies in R' is a subset of the set of his strategies in R. So to prove that s is a Nash equilibrium of R' it suffices to prove that no strategy constituting s is eliminated. Suppose otherwise. Then some s_i is eliminated, so for some $s'_i \in R_i$

$$p_i(s'_i, s''_{-i}) > p_i(s_i, s''_{-i}) \text{ for all } s''_{-i} \in R_{-i}.$$

In particular

$$p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i}),$$

so s is not a Nash equilibrium of R.

(ii) Suppose s is not a Nash equilibrium of R. Then for some $i \in \{1, \ldots, n\}$ strategy s_i is not a best response of player i to s_{-i} in R.

Let $s'_i \in R_i$ be a best response of player i to s_{-i} in R (which exists since R_i is finite). The strategy s'_i is eliminated since s is a Nash equilibrium of R'. So for some $s^*_i \in R_i$

$$p_i(s^*_i, s''_{-i}) > p_i(s'_i, s''_{-i}) \text{ for all } s''_{-i} \in R_{-i}.$$

In particular

$$p_i(s^*_i, s_{-i}) > p_i(s'_i, s_{-i}),$$

which contradicts the choice of s'_i. □

In general an elimination of strictly dominated strategies is not a one step process; it is an iterative procedure. Its use is justified by the assumption of common knowledge of rationality.

Example 7 Consider the following game:

	L	M	R
T	$3,0$	$2,1$	$1,0$
C	$2,1$	$1,1$	$1,0$
B	$0,1$	$0,1$	$0,0$

Note that B is strictly dominated by T and R is strictly dominated by M. By eliminating these two strategies we get:

	L	M
T	$3,0$	$2,1$
C	$2,1$	$1,1$

Now C is strictly dominated by T, so we get:

	L	M
T	$3,0$	$2,1$

In this game L is strictly dominated by M, so we finally get:

	M
T	$2,1$

$\square$

This brings us to the following notion, where given a binary relation $\rightarrow$ we denote by $\rightarrow^*$ its transitive reflexive closure. Consider a strategic game G. Suppose that $G \rightarrow_S^* R$, i.e., R is obtained by an iterated elimination of strictly dominated strategies, in short ***IESDS***, starting with G.

- If for no restriction R' of G, $R \rightarrow_S R'$ holds, we say that R is ***an outcome of IESDS from*** G.
- If each player is left in R with exactly one strategy, we say that G ***is solved by IESDS***.

The following result then clarifies the relation between the IESDS and Nash equilibrium.

Theorem 2 (IESDS) *Suppose that G' is an outcome of IESDS from a strategic game G.*

(i) If s is a Nash equilibrium of G, then it is a Nash equilibrium of G'.

(ii) If G is finite and s is a Nash equilibrium of G', then it is a Nash equilibrium of G.

(iii) If G is finite and solved by IESDS, then the resulting joint strategy is a unique Nash equilibrium.

Proof. By the Strict Elimination Lemma 1. $\square$

Corollary 3 (Strict Dominance) *Consider a strategic game G.*

Suppose that s is a joint strategy such that each s_i is a strictly dominant strategy. Then it is a Nash equilibrium of G. Moreover, if G is finite, then it is a unique Nash equilibrium.

Proof. The first claim is by the definition of Nash equilibrium, while the second one follows by the IESDS Theorem 2(iii). $\square$

Example 8 A nice example of a game that is solved by IESDS is the ***location game***. Assume that that the players are two vendors who simultaneously choose a location. Then the customers choose the closest vendor. The profit for each vendor equals the number of customers it attracted.

To be more specific we assume that the vendors choose a location from the set $\{1, \ldots, n\}$ of natural numbers, viewed as points on a real line, and that at each location there is exactly one customer. For example, for $n = 11$ we have 11 locations:

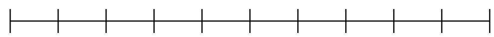

and when the players choose respectively the locations 3 and 8:

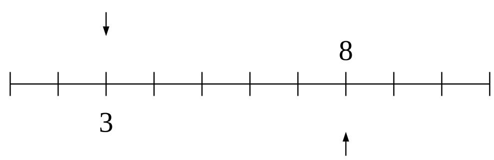

we have $p_1(3, 8) = 5$ and $p_2(3, 8) = 6$. When the vendors 'share' a customer, for instance when they both choose the location 6:

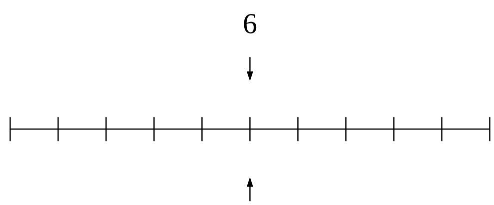

they end up with a fractional payoff, in this case $p_1(6, 6) = 5.5$ and $p_1(6, 6) = 5.5$.

In general, we have the following game:

- each set of strategies consists of the set $\{1, \ldots, n\}$,
- each payoff function p_i is defined by:

$$p_i(s_i, s_{-i}) := \begin{cases} \dfrac{s_i + s_{-i} - 1}{2} & \text{if } s_i < s_{-i} \\ n - \dfrac{s_i + s_{-i} - 1}{2} & \text{if } s_i > s_{-i} \\ \dfrac{n}{2} & \text{if } s_i = s_{-i} \end{cases}$$

It is easy to check that for $n = 2k + 1$ this game is solved by k rounds of IESDS, and that each player is left with the 'middle' strategy k. In each round both 'outer' strategies are eliminated, so first 1 and n, then 2 and $n - 1$, and so on. □

There is one more natural question that we left so far unanswered. Is the outcome of an iterated elimination of strictly dominated strategies unique, or in the game theory parlance: is strict dominance ***order independent***? The answer is positive.

Theorem 4 (Order Independence I) *Given a finite strategic game all iterated eliminations of strictly dominated strategies yield the same outcome.*

Proof. See the Appendix of this Chapter. □

The above result does not hold for infinite strategic games.

Example 9 Consider a game in which the set of strategies for each player is the set of natural numbers. The payoff to each player is the number (strategy) he selected.

Note that in this game every strategy is strictly dominated. Consider now three ways of using IESDS:

- by removing in one step all strategies that are strictly dominated,
- by removing in one step all strategies different from 0 that are strictly dominated,
- by removing in each step exactly one strategy, for instance the least even strategy.

In the first case we obtain the restriction with the empty strategy sets, in the second one we end up with the restriction in which each player has just one strategy, 0, and in the third case we obtain an infinite sequence of reductions. □

The above example also shows that in the limit of an infinite sequence of reductions different outcomes can be reached. So for infinite games the definition of the order independence has to be modified.

The above example also shows that in the Strict Elimination 1(*ii*) and the IESDS Theorem 2(*ii*) and (*iii*) we cannot drop the assumption that the game is finite. Indeed, the above infinite game has no Nash equilibria, while the game in which each player has exactly one strategy has a Nash equilibrium.

Exercise 5

(i) What is the outcome of IESDS in the location game with an even number of locations?

(ii) Modify the location game from Example 8 to a game for three players. Prove that this game has no Nash equilibrium.

(iii) Define a modification of the above game for three players to the case when the set of possible locations (both for the vendors and the customers) forms a circle. Find the set of Nash equilibria. □

Appendix

We provide here the proof of the Order Independence I Theorem 4. Conceptually it is useful to carry out these consideration in a more general setting. We assume an initial strategic game

$$G := (G_1, \ldots, G_n, p_1, \ldots, p_n).$$

By a ***dominance relation*** D we mean a function that assigns to each restriction R of G a subset D_R of $\bigcup_{i=1}^n R_i$. Instead of writing $s_i \in D_R$ we say that s_i ***is D-dominated in*** R.

Given two restrictions R and R' we write $R \rightarrow_D R'$ when $R \neq R', R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\}\ \forall s_i \in R_i \setminus R'_i \ \ s_i \text{ is } D\text{-dominated in } R.$$

Clearly being strictly dominated by another strategy is an example of a dominance relation and $\rightarrow_S$ is an instance of $\rightarrow_D$.

An ***outcome*** of an iteration of $\rightarrow_D$ starting in a game G is a restriction R that can be reached from G using $\rightarrow_D$ in finitely many steps and such that for no R', $R \rightarrow_D R'$ holds.

We call a dominance relation D

- ***order independent*** if for all initial finite games G all iterations of $\rightarrow_D$ starting in G yield the same final outcome,

- ***hereditary*** if for all initial games G, all restrictions R and R' such that $R \rightarrow_D R'$ and a strategy s_i in R'

  s_i is D-dominated in R implies that s_i is D-dominated in R'.

We now establish the following general result.

Theorem 5 *Every hereditary dominance relation is order independent.*

To prove it we introduce the notion of an ***abstract reduction system***. It is simply a pair $(A, \rightarrow)$ where A is a set and $\rightarrow$ is a binary relation on A. Recall that $\rightarrow^*$ denotes the transitive reflexive closure of $\rightarrow$.

- We say that b is a $\rightarrow$***-normal form of*** a if $a \rightarrow^* b$ and no c exists such that $b \rightarrow c$, and omit the reference to $\rightarrow$ if it is clear from the context. If every element of A has a unique normal form, we say that $(A, \rightarrow)$ (or just $\rightarrow$ if A is clear from the context) satisfies the ***unique normal form property***.

- We say that $\rightarrow$ is ***weakly confluent*** if for all $a, b, c \in A$

$$\begin{tikzcd} & a \arrow[dl] \arrow[dr] & \\ b & & c \end{tikzcd}$$

  implies that for some $d \in A$

$$\begin{tikzcd} b \arrow[dr, "*"] & & c \arrow[dl, "*"'] \\ & d & \end{tikzcd}$$

We need the following crucial lemma.

Lemma 6 (Newman) *Consider an abstract reduction system* $(A, \rightarrow)$ *such that*

- *no infinite $\rightarrow$ sequences exist,*
- *$\rightarrow$ is weakly confluent.*

Then $\rightarrow$ satisfies the unique normal form property.

Proof. By the first assumption every element of A has a normal form. To prove uniqueness call an element a *ambiguous* if it has at least two different normal forms. We show that for every ambiguous a some ambiguous b exists such that $a \rightarrow b$. This proves absence of ambiguous elements by the first assumption.

So suppose that some element a has two distinct normal forms n_1 and n_2. Then for some b, c we have $a \rightarrow b \rightarrow^* n_1$ and $a \rightarrow c \rightarrow^* n_2$. By weak confluence some d exists such that $b \rightarrow^* d$ and $c \rightarrow^* d$. Let n_3 be a normal form of d. It is also a normal form of b and of c. Moreover $n_3 \neq n_1$ or $n_3 \neq n_2$. If $n_3 \neq n_1$, then b is ambiguous and $a \rightarrow b$. And if $n_3 \neq n_2$, then c is ambiguous and $a \rightarrow c$. □

Proof of Theorem 5.

Take a hereditary dominance relation D. Consider a restriction R. Suppose that $R \rightarrow_D R'$ for some restriction R'. Let R'' be the restriction of R obtained by removing all strategies that are D-dominated in R.

We have $R'' \subseteq R'$. Assume that $R' \neq R''$. Choose an arbitrary strategy s_i such that $s_i \in R'_i \setminus R''_i$. So s_i is D-dominated in R. By the hereditarity of D, s_i is also D-dominated in R'. This shows that $R' \rightarrow_D R''$.

So we proved that either $R' = R''$ or $R' \rightarrow_D R''$, i.e., that $R' \rightarrow_D^* R''$. This implies that $\rightarrow_D$ is weakly confluent. It suffices now to apply Newman's Lemma 6. □

To apply this result to strict dominance we establish the following fact.

Lemma 7 (Hereditarity I) *The relation of being strictly dominated is hereditary on the set of restrictions of a given finite game.*

Proof. Suppose a strategy $s_i \in R'_i$ is strictly dominated in R and $R \rightarrow_S R'$. The initial game is finite, so there exists in R_i a strategy s'_i that strictly dominates s_i in R and is not strictly dominated in R. Then s'_i is not eliminated in the step $R \rightarrow_S R'$ and hence is a strategy in R'_i. But $R' \subseteq R$, so s'_i also strictly dominates s_i in R'. □

The promised proof is now immediate.

Proof of the Order Independence I Theorem 4.
By Theorem 5 and the Hereditarity I Lemma 7. □

Chapter 4

Weak Dominance and Never Best Responses

Let us return now to our analysis of an arbitrary strategic game $G := (S_1, \ldots, S_n,\ p_1, \ldots, p_n)$. Let s_i, s'_i be strategies of player i. We say that s_i ***weakly dominates*** s'_i (or equivalently, that s'_i is ***weakly dominated by*** s_i) if

$$\forall s_{-i} \in S_{-i}\ p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}) \text{ and } \exists s_{-i} \in S_{-i}\ p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

Further, we say that s_i is ***weakly dominant*** if it weakly dominates all other strategies of player i.

4.1 Elimination of weakly dominated strategies

Analogous considerations to the ones concerning strict dominance can be carried out for the elimination of weakly dominated strategies. To this end we consider the reduction relation $\rightarrow_W$ on the restrictions of G, defined by

$$R \rightarrow_W R'$$

when $R \neq R'$, $R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\}\ \forall s_i \in R_i \setminus R'_i\ \exists s'_i \in R_i\ s_i \text{ is weakly dominated in } R \text{ by } s'_i.$$

Below we abbreviate iterated elimination of weakly dominated strategies to ***IEWDS***.

However, in the case of IEWDS some complications arise. To illustrate them consider the following game that results from equipping each player in the Matching Pennies game with a third strategy E (for Edge):

	H	T	E
H	$1, -1$	$-1, \ \ 1$	$-1, -1$
T	$-1, \ \ 1$	$1, -1$	$-1, -1$
E	$-1, -1$	$-1, -1$	$-1, -1$

Note that

- (E, E) is its only Nash equilibrium,
- for each player E is the only strategy that is weakly dominated.

Any form of elimination of these two E strategies, simultaneous or iterated, yields the same outcome, namely the Matching Pennies game, that, as we have already noticed, has no Nash equilibrium. So during this eliminating process we 'lost' the only Nash equilibrium. In other words, part (i) of the IESDS Theorem 2 does not hold when reformulated for weak dominance.

On the other hand, some partial results are still valid here. As before we prove first a lemma that clarifies the situation.

Lemma 8 (Weak Elimination) *Given a finite strategic game G consider two restrictions R and R' of G such that $R \rightarrow_W R'$. Then if s is a Nash equilibrium of R', then it is a Nash equilibrium of R.*

Proof. Suppose s is a Nash equilibrium of R' but not a Nash equilibrium of R. Then for some $i \in \{1, \ldots, n\}$ the set

$$A := \{s'_i \in R_i \mid p_i(s'_i, s_{-i}) > p_i(s)\}$$

is non-empty.

Weak dominance is a strict partial ordering (i.e. an irreflexive transitive relation) and A is finite, so some strategy s'_i in A is not weakly dominated in R by any strategy in A. But each strategy in A is eliminated in the reduction

$R \rightarrow_W R'$ since s is a Nash equilibrium of R'. So some strategy $s_i^* \in R_i$ weakly dominates s_i' in R. Consequently

$$p_i(s_i^*, s_{-i}) \geq p_i(s_i', s_{-i})$$

and as a result $s_i^* \in A$. But this contradicts the choice of s_i'. $\Box$

This brings us directly to the following result.

Theorem 9 (IEWDS) *Suppose that G is a finite strategic game.*

(i) If G' is an outcome of IEWDS from G and s is a Nash equilibrium of G', then s is a Nash equilibrium of G.

(ii) If G is solved by IEWDS, then the resulting joint strategy is a Nash equilibrium of G.

Proof. By the Weak Elimination Lemma 8. $\Box$

Corollary 10 (Weak Dominance) *Consider a finite strategic game G.*

Suppose that s is a joint strategy such that each s_i is a weakly dominant strategy. Then it is a Nash equilibrium of G.

Proof. By the IEWDS Theorem 9(*ii*). $\Box$

Note that in contrast to the Strict Dominance Corollary 3 we do not claim here that s is a unique Nash equilibrium of G. In fact, such a stronger claim does not hold. Indeed, consider the game

	L	R
T	$1, 1$	$1, 1$
B	$1, 1$	$0, 0$

Here T is a weakly dominant strategy for the player 1, L is a weakly dominant strategy for player 2 and, as prescribed by the above Note, (T, L), is a Nash equilibrium. However, this game has two other Nash equilibria, (T, R) and (B, L).

Example 10 Let us return to the beauty contest game introduced in Example 2 of Chapter 1. One can check that this game is solved by IEWDS and results in the joint strategy $(1, \ldots, 1)$. Hence, we can conclude by the IEWDS Theorem 9 this joint strategy is a (not necessarily unique; we shall return to this question in a later chapter) Nash equilibrium. $\Box$

Note that in contrast to the IESDS Theorem 2 we do not claim in part (ii) of the IEWDS Theorem 9 that the resulting joint strategy is a *unique* Nash equilibrium. In fact, such a stronger claim does not hold. Further, in contrast to strict dominance, an iterated elimination of weakly dominated strategies can yield several outcomes.

The following example reveals even more peculiarities of this procedure.

Example 11 Consider the following game:

	L	M	R
T	$0,1$	$1,0$	$0,0$
B	$0,0$	$0,0$	$1,0$

It has three Nash equlibria, (T, L), (B, L) and (B, R). This game can be solved by IEWDS but only if in the first round we do not eliminate all weakly dominated strategies, which are M and R. If we eliminate only R, then we reach the game

	L	M
T	$0,1$	$1,0$
B	$0,0$	$0,0$

that is solved by IEWDS by eliminating B and M. This yields

	L
T	$0,1$

So not only IEWDS is not order independent; in some games it is advantageous *not* to proceed with the deletion of the weakly dominated strategies 'at full speed'. One can also check that the second Nash equilibrium, (B, L), can be found using IEWDS, as well, but not the third one, (B, R). $\square$

It is instructive to see where the proof of order independence given in the Appendix of the previous chapter breaks down in the case of weak dominance. This proof crucially relied on the fact that the relation of being strictly dominated is hereditary. In contrast, the relation of being weakly dominated is not hereditary.

To summarize, the iterated elimination of weakly dominated strategies

- can lead to a deletion of Nash equilibria,

- does not need to yield a unique outcome,
- can be too restrictive if we stipulate that in each round all weakly dominated strategies are eliminated.

Finally, note that the above IEWDS Theorem 9 does not hold for infinite games. Indeed, Example 9 applies here, as well.

4.2 Elimination of never best responses

Iterated elimination of strictly or weakly dominated strategies allow us to solve various games. However, several games cannot be solved using them.

For example, consider the following game:

	X	Y
A	$2,1$	$0,0$
B	$0,1$	$2,0$
C	$1,1$	$1,2$

Here no strategy is strictly or weakly dominated. On the other hand C is a ***never best response***, that is, it is not a best response to any strategy of the opponent. Indeed, A is a unique best response to X and B is a unique best response to Y. Clearly, the above game is solved by an iterated elimination of never best responses. So this procedure can be stronger than IESDS and IEWDS.

Formally, we introduce the following reduction notion between the restrictions R and R' of a given strategic game G:

$$R \rightarrow_N R'$$

when $R \neq R'$, $R' \subseteq R$ and

$$\forall i \in \{1, \ldots, n\} \; \forall s_i \in R_i \setminus R'_i \; \neg \exists s_{-i} \in R_{-i} \; s_i \text{ is a best response to } s_{-i} \text{ in } R.$$

That is, $R \rightarrow_N R'$ when R' results from R by removing from it some strategies that are never best responses. Note that in contrast to strict and weak dominance there is now no 'witness' strategy that acounts for a removal of a strategy.

We now focus on the iterated elimination of never best responses, in short ***IENBR***, obtained by using the $\rightarrow_N^*$ relation. The following counterpart of the IESDS Theorem 2 holds.

Theorem 11 (IENBR) *Suppose that G' is an outcome of IENBR from a strategic game G.*

(i) If s is a Nash equilibrium of G, then it is a Nash equilibrium of G'.

(ii) If G is finite and s is a Nash equilibrium of G', then it is a Nash equilibrium of G.

(iii) If G is finite and solved by IENBR, then the resulting joint strategy is a unique Nash equilibrium.

Proof. Analogous to the proof of the IESDS Theorem 2 and omitted. $\Box$

Further, we have the following analogue of the Hereditarity I Lemma 7.

Lemma 12 (Hereditarity II) *The relation of never being a best response is hereditary on the set of restrictions of a given finite game.*

Proof. Suppose a strategy $s_i \in R'_i$ is a never best response in R and $R \rightarrow_N R'$. Assume by contradiction that for some $s_{-i} \in R'_{-i}$, s_i is a best response to s_{-i} in R', i.e.,

$$\forall s'_i \in R'_i \; p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

The initial game is finite, so there exists a best response s'_i to s_{-i} in R. Then s'_i is not eliminated in the step $R \rightarrow_N R'$ and hence is a strategy in R'_i. But s_i is not a best response to s_{-i} in R, so

$$p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i}),$$

so we reached a contradiction. $\Box$

This leads us to the following analogue of the Order Independence I Theorem 4.

Theorem 13 (Order Independence II) *Given a finite strategic game all iterated eliminations of never best responses yield the same outcome.*

Proof. By Theorem 5 and the Hereditarity II Lemma 12. $\Box$

In the case of infinite games we encounter the same problems as in the case of IESDS. Indeed, Example 9 readily applies to IENBR, as well, since in this game no strategy is a best response. In particular, this example shows that if we solve an infinite game by IENBR we cannot claim that we obtained a Nash equilibrium. Still, IENBR can be useful in such cases.

Example 12 Consider the following infinite variant of the location game considered in Example 8. We assume that the players choose their strategies from the open interval $(0, 100)$ and that at each real in $(0, 100)$ there resides one customer. We have then the following payoffs that correspond to the intuition that the customers choose the closest vendor:

$$p_i(s_i, s_{-i}) := \begin{cases} \dfrac{s_i + s_{-i}}{2} & \text{if } s_i < s_{-i} \\ 100 - \dfrac{s_i + s_{-i}}{2} & \text{if } s_i > s_{-i} \\ 50 & \text{if } s_i = s_{-i} \end{cases}$$

It is easy to check that in this game no strategy strictly or weakly dominates another one. On the other hand each strategy 50 is a best response (namely to strategy 50 of the opponent) and no other strategies are best responses. So this game is solved by IENBR, in one step.

We cannot claim automatically that the resulting joint strategy $(50, 50)$ is a Nash equilibrium, but it is clearly so since each strategy 50 is a best response to the 'other' strategy 50. Moreover, by the IENBR Theorem 11(i) we know that this is a unique Nash equilibrium. □

Exercise 6 Show that the beauty contest game from Example 2 is indeed solved by IEWDS. □

Exercise 7 Show that in the location game from Example 12 indeed no strategy is strictly or weakly dominated. □

Exercise 8 Given a game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ we say that that a strategy s_i of player i is ***dominant*** if for all strategies s'_i of player i

$$p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}).$$

Suppose that s is a joint strategy such that each s_i is a dominant strategy. Prove that it is a Nash equilibrium of G. □

Chapter 5

Regret Minimization and Security Strategies

Until now we implicitly adopted a view that a Nash equilibrium is a desirable outcome of a strategic game. In this chapter we consider two alternative views that help us to understand reasoning of players who either want to avoid costly 'mistakes' or 'fear' a bad outcome. Both concepts can be rigorously formalized.

5.1 Regret minimization

Consider the following game:

	L	R
T	$100, 100$	$0, \quad 0$
B	$0, \quad 0$	$1, \quad 1$

This is an example of a coordination problem, in which there are two satisfactory outcomes (read Nash equilibria), (T, L) and (B, R), of which one is obviously better for both players. In this game no strategy strictly or weakly dominates the other and each strategy is a best response to some other strategy. So using the concepts we introduced so far we cannot explain how come that rational players would end up choosing the Nash equilibrium (T, L). In this section we explain how this choice can be justified using the concept of ***regret minimization***.

With each finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ we first associate a ***regret-recording game*** $G := (S_1, \ldots, S_n, r_1, \ldots, r_n)$ in which each payoff function r_i is defined by

$$r_i(s_i, s_{-i}) := p_i(s_i^*, s_{-i}) - p_i(s_i, s_{-i}),$$

where s_i^* is player's i best response to s_{-i}. We call then $r_i(s_i, s_{-i})$ player's i ***regret of choosing*** s_i ***against*** s_{-i}. Note that by definition for all s we have $r_i(s) \geq 0$.

For example, for the above game the corresponding regret-recording game is

	L	R
T	$0, \quad 0$	$1, 100$
B	$100, \quad 1$	$0, \quad 0$

Indeed, $r_1(B, L) := p_1(T, L) - p_1(B, L) = 100$, and similarly for the other seven entries.

Let now

$$regret_i(s_i) := \max_{s_{-i} \in S_{-i}} r_i(s_i, s_{-i}).$$

So $regret_i(s_i)$ is the maximal regret player i can have from choosing s_i. We call then any strategy s_i^* for which the function $regret_i$ attains the minimum, i.e., one such that $regret_i(s_i^*) = \min_{s_i \in S_i} regret_i(s_i)$, a ***regret minimization strategy*** for player i.

In other words, s_i^* is a regret minimization strategy for player i if

$$\max_{s_{-i} \in S_{-i}} r_i(s_i^*, s_{-i}) = \min_{s_i \in S_i} \max_{s_{-i} \in S_{-i}} r_i(s_i, s_{-i}).$$

The following intuition is helpful here. Suppose the opponents of player i are able to perfectly anticipate which strategy player i is about to play (for example by being informed through a third party what strategy player i has just selected and is about to play). Suppose further that they aim at inflicting at player i the maximum damage in the form of maximal regret and that player i is aware of these circumstances. Then to miminize his regret player i should select a regret minimization strategy. We could say that a regret minimization strategy will be chosen by a player who wants to avoid making a costly 'mistake', where by a mistake we mean a choice of a strategy that is not a best response to the joint strategy of the opponents.

To clarify this notion let us return to our example of the coordination game. To visualize the outcomes of the functions $regret_1$ and $regret_2$ we put the results in an additional row and column:

	L	R	$regret_1$
T	$0, \quad 0$	$1, 100$	1
B	$100, \quad 1$	$0, \quad 0$	100
$regret_2$	1	100	

So T is the minimum of $regret_1$ and L is the minimum of $regret_2$. Hence (T, L) is the unique pair of regret minimization strategies. This shows that using the concept of regret minimization we succeeded to single out the preferred Nash equilibrium in the considered coordination game.

It is important to note that the concept of regret minimization does not allow us to solve all coordination problems. For example, it does not help us in selecting a Nash equilibrium in symmetric situations, for instance in the game

	L	R
T	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Indeed, in this case the regret of each strategy is 1, so regret minimization does not allow us to distinguish between the strategies. Analogous considerations hold for the Battle of Sexes game from Chapter 1.

Regret minimization is based on different intuitions than strict and weak dominance. As a result these notions are incomparable. In general, only the following limited observation holds. Recall that the notion of a dominant strategy was introduced in Exercise 8 on page 32.

Note 14 (Regret Minimization) *Consider a finite game. Every dominant strategy is a regret minimization strategy.*

Proof. Fix a finite game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. Note that each dominant strategy of player i is a best response to each $s_{-i} \in S_{-i}$. So by the definition of the regret-recording game for all $s_{-i} \in S_{-i}$ we have $r_i(s_i, s_{-i}) = 0$. This shows that s_i is a regret minimization strategy for player i, since for all joint strategies s we have $r_i(s) \geq 0$. $\square$

The process of removing strategies that do not achieve regret minimization can be iterated. We call this process the ***iterated regret minimization***. The example of the coordination game we analyzed shows that the process of regret minimization may yield to a loss of some Nash equilibria. In fact, as we shall see in a moment, during this process all Nash equilibria can be lost. On the other hand, as recently suggested by J. Halpern and R. Pass, in some games the iterated regret minimization yields a more intuitive outcome. As an example let us return to the Traveller's Dilemma game considered in Example 1.

Example 13 (Traveller's dilemma revisited)

Let us first determine in this game the regret minimization strategies for each player. Take a joint strategy s.

Case 1. $s_{-i} = 2$.

Then player's i regret of choosing s_i against s_{-i} is 0 if $s_i = s_{-i}$ and 2 if $s_i > s_{-i}$, so it is at most 2.

Case 2. $s_{-i} > 2$.

If $s_{-i} < s_i$, then $p_i(s) = s_{-i} - 2$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case 3.

If $s_{-i} = s_i$, then $p_i(s) = s_{-i}$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case 1.

Finally, if $s_{-i} > s_i$, then $p_i(s) = s_i + 2$, while the best response to s_{-i}, namely $s_{-i} - 1$, yields the payoff $s_{-i} + 1$. So player's i regret of choosing s_i against s_{-i} is in this case $s_{-i} - s_i - 1$.

To summarize, we have

$$regret_i(s_i) = \max(3, \max_{s_{-i} \in S_{-i}} s_{-i} - s_i - 1) = \max(3, 99 - s_i).$$

So the minimal regret is achieved when $99 - s_i \leq 3$, i.e., when the strategy s_i is in the interval $[96, 100]$. Hence removing all strategies that do not achieve regret minimization yields a game in which each player has the strategies in the interval $[96, 100]$. In particular, we 'lost' in this way the unique Nash equilibrium of this game, (2,2).

We now repeat this elimination procedure. To compute the outcome we consider again two, though now different, cases.

Case 1. $s_i = 97$.

The following table then summarizes player's i regret of choosing s_i against a strategy s_{-i} of player i:

strategy of player $-i$	best response of player i	regret of player i
96	96	2
97	96	1
98	97	0
99	98	1
100	99	2

Case 2. $s_i \neq 97$.

The following table then summarizes player's i regret of choosing s_i, where for each strategy of player i we list a strategy of player $-i$ for which player's i regret is maximal:

strategy of player i	relevant strategy of player $-i$	regret of player i
96	100	3
98	97	3
99	98	3
100	99	3

So each strategy of player i different from 97 has regret 3, while 97 has regret 2. This means that the second round of elimination of the strategies that do not achieve regret minimization yields a game in which each player has just one strategy, namely 97. □

Recall again that the unique Nash equilibrium in the Traveller's Dilemma game is (2,2). So the iterated regret minimization yields here a radically different outcome than the analysis based on Nash equilibria. Interestingly, this outcome, (97,97), has been confirmed by empirical studies.

We conclude this section by showing that iterated regret minimization is not order independent. To this end consider the following game:

	L	R
T	$2, 1$	$0, 3$
B	$0, 2$	$1, 1$

The corresponding regret-recording game, together with the recording of the outcomes of the functions $regret_1$ and $regret_2$ is as follows:

	L	R	$regret_1$
T	$0, 2$	$1, 0$	1
B	$2, 0$	$0, 1$	2
$regret_2$	2	1	

This shows that (T, R) is the unique pair of regret minimization strategies in the original game. So by removing from the original game the strategies B and L that do not achieve regret minimization we reduce it to

	R
T	$0, 3$

On the other hand, if we initially only remove strategy L, then we obtain the game

	R
T	$0, 3$
B	$1, 1$

Now the only strategy that does not achieve regret minimization is T. By removing it we obtain the game

	R
B	$1, 1$

5.2 Security strategies

Consider the following game:

	L	R
T	$0, \quad 0$	$101, \quad 1$
B	$1, 101$	$100, 100$

This is an extreme form of a ***Chicken game***, sometimes also called a ***Hawk-Dove game*** or a ***Snowdrift game***.

The game of Chicken models two drivers driving at each other on a narrow road. If neither driver swerves ('chickens'), the result is a crash. The best option for each driver is to stay straight while the other swerves. This yields a situation where each driver, in attempting to realize his the best outcome, risks a crash.

The description of this game as a snowdrift game stresses advantages of a cooperation. The game involves two drivers who are trapped on opposite sides of a snowdrift. Each has the option of staying in the car or shoveling snow to clear a path. Letting the other driver do all the work is the best option, but being exploited by shoveling while the other driver sits in the car is still better than doing nothing.

Note that this game has two Nash equilibria, (T, R) and (B, L). However, there seems to be no reason in selecting any Nash equilibrium as each Nash equlibrium is grossly unfair to the player who will receive only 1.

In contrast, (B, R), which is not a Nash equilibrium, looks like a most reasonable outcome. Each player receives in it a payoff close to the one he receives in the Nash equilibrium of his preference. Also, why should a player risk the payoff 0 in his attempt to secure the payoff 101 that is only a fraction bigger than his payoff 100 in (B, R)?

Note that in this game no strategy strictly or weakly dominates the other and each strategy is a best response to some other strategy. So these concepts are useless in analyzing this game. Moreover, the regret minimization for each strategy is 1. So this concept is of no use here either.

We now introduce the concept of a ***security strategy*** that allows us to single out the joint strategy (B, R) as the most reasonable outcome for both players.

Fix a, not necessarily finite, strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$. Player i, when considering which strategy s_i to select, has to take into account which strategies his opponents will choose. A 'worst case scenario' for player i is that, given his choice of s_i, his opponents choose a joint strategy for which player's i payoff is the lowest[1]. For each strategy s_i of player i once this lowest payoff can be identified a strategy can be selected that leads to a 'minimum damage'.

[1]We assume here that such s_i exists.

To formalize this concept for each $i \in \{1, \ldots, n\}$ we consider the function[2]

$$f_i : S_i \to \mathcal{R}$$

defined by

$$f_i(s_i) := \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

We call any strategy s_i^* for which the function f_i attains the maximum, i.e., one such that $f_i(s_i^*) = \max_{s_i \in S_i} f_i(s_i)$, a ***security strategy*** or a ***maxminimizer*** for player i. We denote this maximum, so

$$\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}),$$

by $maxmin_i$ and call it the ***security payoff*** of player i.

In other words, s_i^* is a security strategy for player i if

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i.$$

Note that $f_i(s_i)$ is the minimum payoff player i is guaranteed to secure for himself when he selects strategy s_i. In turn, the security payoff $maxmin_i$ of player i is the minimum payoff he is guaranteed to secure for himself in general. To achieve at least this payoff he just needs to select any security strategy.

The following intuition is helpful here. Suppose the opponents of player i are able to perfectly anticipate which strategy player i is about to play. Suppose further that they aim at inflicting at player i the maximum damage (in the form of the lowest payoff) and that player i is aware of these circumstances. Then player i should select a strategy that causes the minimum damage for him. Such a strategy is exactly a security strategy and it guarantees him at least the $maxmin_i$ payoff. We could say that a security strategy will be chosen by a 'pessimist' player, i.e., one who fears the worst outcome for himself.

To clarify this notion let us return to our example of the chicken game. Clearly, both B and R are the only security strategies in this game. Indeed, we have $f_1(T) = f_2(L) = 0$ and $f_1(B) = f_2(R) = 1$. So we succeeded to

[2] In what follows we assume that all considered minima and maxima always exist. This assumption is obviously satisfied in finite games. In a later chapter we shall discuss a natural class of infinite games for which this assumption is satisfied, as well.

single out in this game the outcome (B, R) using the concept of a security strategy.

The following counterpart of the Regret Minimization Note 14 holds.

Note 15 (Security) *Consider a finite game. Every dominant strategy is a security strategy.*

Proof. Fix a game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ and suppose that s_i^* is a dominant strategy of player i. For all joint strategies s

$$p_i(s_i^*, s_{-i}) \geq p_i(s_i, s_{-i}),$$

so for all strategies s_i of player i

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) \geq \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

Hence

$$\min_{s_{-i} \in S_{-i}} p_i(s_i^*, s_{-i}) \geq \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

This concludes the proof. □

Next, we introduce a dual notion to the security payoff $maxmin_i$. It is not needed for the analysis of security strategies but it will turn out to be relevant in a later chapter.

With each $i \in \{1, \ldots, n\}$ we consider the function

$$F_i : S_{-i} \to \mathcal{R}$$

defined by

$$F_i(s_{-i}) := \max_{s_i \in S_i} p_i(s_i, s_{-i}).$$

Then we denote the value $\min_{s_{-i} \in S_{-i}} F_i(s_{-i})$, i.e.,

$$\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} p_i(s_i, s_{-i}),$$

by $minmax_i$.

The following intuition is helpful here. Suppose that now player i is able to perfectly anticipate which strategies his opponents are about to play. Using this information player i can compute the minimum payoff he is guaranteed to achieve in such circumstances: it is $minmax_i$. This lowest payoff for player

i can be enforced by his opponents if they choose any joint strategy s^*_{-i} for which the function F_i attains the minimum, i.e., one such that $F_i(s^*_{-i}) = \min_{s_{-i} \in S_{-i}} F_i(s_{-i})$.

To clarify the notions of $maxmin_i$ and $minmax_i$ consider an example.

Example 14 Consider the following two-player game:

	L	M	R
T	$3, -$	$4, -$	$5, -$
B	$6, -$	$2, -$	$1, -$

where we omit the payoffs of the second, i.e., column, player.

To visualize the outcomes of the functions f_1 and F_1 we put the results in an additional row and column:

	L	M	R	f_1
T	$3, -$	$4, -$	$5, -$	3
B	$6, -$	$2, -$	$1, -$	1
F_1	6	4	5	

That is, in the f_1 column we list for each row its minimum and in the F_1 row we list for each column its maximum.

Since $f_1(T) = 3$ and $f_1(B) = 1$ we conclude that $maxmin_1 = 3$. So the security payoff of the row player is 3 and T is a unique security strategy of the row player. In other words, the row player can secure for himself at least the payment 3 and achieves this by choosing strategy T.

Next, since $F_1(L) = 6, F_1(M) = 4$ and $F_1(R) = 5$ we get $minmax_1 = 4$. In other words, if the row player knows which strategy the column player is to play, he can secure for himself at least the payment 4.

Indeed,

- if the row player knows that the column player is to play L, then he should play B (and secure the payoff 6),
- if the row player knows that the column player is to play M, then he should play T (and secure the payoff 4),
- if the row player knows that the column player is to play R, then he should play T (and secure the payoff 5). $\Box$

In the above example $maxmin_1 < minmax_1$. In general the following observation holds. From now on, to simplify the notation we assume that s_i and s_{-i} range over, respectively, S_i and S_{-i}.

Lemma 16 (Lower Bound)

(i) For all $i \in \{1, \ldots, n\}$ we have $maxmin_i \leq minmax_i$.

(ii) If s is a Nash equilibrium of G, then for all $i \in \{1, \ldots, n\}$ we have $minmax_i \leq p_i(s)$.

Item (i) formalizes the intuition that one can take a better decision when more information is available (in this case about which strategies the opponents are about to play). Item (ii) provides a lower bound on the payoff in each Nash equilibrium, which explains the name of the lemma.

Proof.
(i) Fix i. Let s_i^* be such that $\min_{s_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i$ and s_{-i}^* such that $\max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i$. We have then the following string of equalities and inequalities:

$$maxmin_i = min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq p_i(s_i^*, s_{-i}^*) \leq max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i.$$

(ii) Fix i. For each Nash equilibrium (s_i^*, s_{-i}^*) of G we have

$$min_{s_{-i}} max_{s_i} p_i(s_i, s_{-i}) \leq max_{s_i} p_i(s_i, s_{-i}^*) = p_i(s_i^*, s_{-i}^*).$$

$\square$

To clarify the difference between the regret minimization and security strategies consider the following variant of a coordination game:

	L	R
T	$100, 100$	$0, \quad 0$
B	$1, \quad 1$	$2, \quad 2$

It is easy to check that players who select the regret minimization strategies will choose the strategies T and L which yields the payoff 100 to each of them. In contrast, players who select the security strategies will choose B and L and will receive only 1 each.

Next, consider the following game:

	L	M	R
T	$5,\ \ 5$	$0,\ \ 0$	$97,\ \ 1$
B	$1,\ \ 0$	$1,\ \ 0$	$100, 100$

Here the security strategies are B and R and their choice by the players yields the payoff 100 to each of them. In contrast, the regret minimization strategies are T (with the regret 3) and R (with the regret 4) and their choice by the players yields them the respective payoffs 97 and 1.

So the outcomes of selecting regret minimization strategies and of security strategies are incomparable.

Finally, note that in general there is no relation between the equalities of the $maxmin_i = minmax_i$ and an existence of a Nash equilibrium. To see this let us fill in the game considered in Example 14 the payoffs for the column player as follows:

	L	M	R
T	$3, 1$	$4, 0$	$5, 1$
B	$6, 1$	$2, 0$	$1, 1$

We already noted that $maxmin_1 < minmax_1$ holds here. However, this game has two Nash equilibria, (T, R) and (B, L).

Further, the following game

	L	M	R
T	$3, 1$	$3, 0$	$5, 0$
B	$6, 0$	$2, 1$	$1, 0$

has no Nash equilibrium and yet for $i = 1, 2$ we have $maxmin_i = minmax_i$.

In a later chapter we shall discuss a class of two-player games for which there is a close relation between the existence of a Nash equilibrium and the equalities $maxmin_i = minmax_i$.

Chapter 6

Strictly Competitive Games

In this chapter we discuss a special class of two-player games for which stronger results concerning Nash equilibria can be established. To study them we shall crucially rely on the notions introduced in Section 5.2, namely security strategies and $maxmin_i$ and $minmax_i$.

More specifically, we introduce a natural class of two-player games for which the equalities between the $maxmin_i$ and $minmax_i$ values for $i = 1, 2$ constitute a necessary and sufficient condition for the existence of a Nash equilibrium. In these games any Nash equilibrium consists of a pair of security strategies.

A ***strictly competitive game*** is a two-player strategic game (S_1, S_2, p_1, p_2) in which for $i = 1, 2$ and any two joint strategies s and s'

$$p_i(s) \geq p_i(s') \text{ iff } p_{-i}(s) \leq p_{-i}(s').$$

That is, a joint strategy that is better for one player is worse for the other player. This formalizes the intuition that the interests of both players are diametrically opposed and explains the terminology.

By negating both sides of the above equivalence we get

$$p_i(s) < p_i(s') \text{ iff } p_{-i}(s) > p_{-i}(s').$$

So an alternative way of defining a strictly competitive game is by stating that this is a two-player game in which every joint strategy is a Pareto efficient outcome.

To illustrate this concept let us fill in the game considered in Example 14 the payoffs for the column player in such a way that the game becomes strictly competitive:

	L	M	R
T	$3,4$	$4,3$	$5,2$
B	$6,0$	$2,5$	$1,6$

Canonic examples of strictly competitive games are ***zero-sum games***. These are two-player games in which for each joint strategy s we have

$$p_1(s) + p_2(s) = 0.$$

So a zero-sum game is an extreme form of a strictly competitive game in which whatever one player 'wins', the other one 'loses'. A simple example is the Matching Pennies game from Chapter 1.

Another well-known zero-sum game is the Rock, Paper, Scissors game. In this game, often played by children, both players simultaneously make a sign with a hand that identifies one of these three objects. If both players make the same sign, the game is a draw. Otherwise one player wins, say, 1 Euro from the other player according to the following rules:

- the rock defeats (breaks) scissors,
- scissors defeat (cut) the paper,
- the paper defeats (wraps) the rock.

Since in a zero-sum game the payoff for the second player is just the negative of the payoff for the first player, each zero-sum game can be represented in a simplified form, called ***reward matrix***. It is simply the matrix that represents only the payoffs for the first player. So the reward matrix for the Rock, Paper, Scissors game looks as follows:

	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

For the strictly competitive games, so a fortiori the zero-sum games, the following counterpart of the Lower Bound Lemma 16 holds.

Lemma 17 (Upper Bound) *Consider a strictly competitive game* $G := (S_1, S_2, p_1, p_2)$. *If* (s_1^*, s_2^*) *is a Nash equilibrium of* G*, then for* $i = 1, 2$

(i) $p_i(s_i^*, s_{-i}^*) \leq \min_{s_{-i}} p_i(s_i^*, s_{-i})$,

(ii) $p_i(s_i^*, s_{-i}^*) \leq maxmin_i$.

Both items provide an upper bound on the payoff in each Nash equilibrium, which explains the name of the lemma.

Proof.

(i) Fix i. Suppose that (s_i^*, s_{-i}^*) is a Nash equilibrium of G. Fix s_{-i}. By the definition of Nash equilibrium

$$p_{-i}(s_i^*, s_{-i}^*) \geq p_{-i}(s_i^*, s_{-i}),$$

so, since G is strictly competitive,

$$p_i(s_i^*, s_{-i}^*) \leq p_i(s_i^*, s_{-i}).$$

But s_{-i} was arbitrary, so

$$p_i(s_i^*, s_{-i}^*) \leq \min_{s_{-i}} p_i(s_i^*, s_{-i}).$$

(ii) By definition

$$\min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq \max_{s_i} \min_{s_{-i}} p_i(s_i, s_{-i}),$$

so by (i)

$$p_i(s_i^*, s_{-i}^*) \leq \max_{s_i} \min_{s_{-i}} p_i(s_i, s_{-i}).$$

□

Combining the Lower Bound Lemma 16 and the Upper Bound Lemma 17 we can draw the following conclusions about strictly competitive games.

Theorem 18 (Strictly Competitive Games) *Consider a strictly competitive game G.*

(i) If for $i = 1, 2$ we have $maxmin_i = minmax_i$, then G has a Nash equilibrium.

(ii) If G has a Nash equilibrium, then for $i = 1, 2$ we have $maxmin_i = minmax_i$.

(iii) All Nash equilibria of G yield the same payoff, namely $maxmin_i$ for player i.

(iv) All Nash equilibria of G are of the form (s_1^*, s_2^*) where each s_i^* is a security strategy for player i.

Proof. Suppose $G = (S_1, S_2, p_1, p_2)$.
(i) Fix i. Let s_i^* be a security strategy for player i, i.e., such that $\min_{s_{-i}} p_i(s_i^*, s_{-i}) = maxmin_i$, and let s_{-i}^* be such that $\max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i$. We show that (s_i^*, s_{-i}^*) is a Nash equilibrium of G.

We already noted in the proof of the Lower Bound Lemma 16(i) that

$$maxmin_i = \min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq p_i(s_i^*, s_{-i}^*) \leq \max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i.$$

But now $maxmin_i = minmax_i$, so all these values are equal. In particular

$$p_i(s_i^*, s_{-i}^*) = \max_{s_i} p_i(s_i, s_{-i}^*) \tag{6.1}$$

and

$$p_i(s_i^*, s_{-i}^*) = \min_{s_{-i}} p_i(s_i^*, s_{-i}).$$

Fix now s_{-i}. By the last equality

$$p_i(s_i^*, s_{-i}^*) \leq p_i(s_i^*, s_{-i}),$$

so, since G is strictly competitive,

$$p_{-i}(s_i^*, s_{-i}^*) \geq p_{-i}(s_i^*, s_{-i}).$$

But s_{-i} was arbitrary, so

$$p_{-i}(s_i^*, s_{-i}^*) = \max_{s_{-i}} p_{-i}(s_i^*, s_{-i}). \tag{6.2}$$

Now (6.1) and (6.2) mean that indeed (s_i^*, s_{-i}^*) is a Nash equilibrium of G.

(ii) and (iii) If s is a Nash equilibrium of G, by the Lower Bound Lemma 16(i) and (ii) and the Upper Bound Lemma 17(ii) we have for $i = 1, 2$

$$maxmin_i \leq minmax_i \leq p_i(s) \leq maxmin_i.$$

So all these values are equal.

(iv) Fix i. Take a Nash equilibrium (s_i^*, s_{-i}^*) of G. We always have

$$\min_{s_{-i}} p_i(s_i^*, s_{-i}) \leq p_i(s_i^*, s_{-i}^*)$$

and by the Upper Bound Lemma 17(i) we also have

$$p_i(s_i^*, s_{-i}^*) \leq \min_{s_{-i}} p_i(s_i^*, s_{-i}).$$

So

$$\min_{s_{-i}} p_i(s_i^*, s_{-i}) = p_i(s_i^*, s_{-i}^*) = maxmin_i,$$

where the last equality holds by (iii). So s_i^* is a security strategy for player i. $\Box$

Combining (i) and (ii) we see that a strictly competitive game has a Nash equilibrium iff for $i = 1, 2$ we have $maxmin_i = minmax_i$. So in a strictly competitive game each player can determine whether a Nash equilibrium exists without knowing the payoff of the other player. All what he needs to know is that the game is strictly competitive. Indeed, each player i then just needs to check whether his $maxmin_i$ and $minmax_i$ values are equal.

Morever, by (iv), each player can select on his own a strategy that forms a part of a Nash equilibrium: it is simply any of his security strategies.

6.1 Zero-sum games

Let us focus now on the special case of zero-sum games. We first show that for zero-sum games the $maxmin_i$ and $minmax_i$ values for one player can be directly computed from the corresponding values for the other player.

Theorem 19 (Zero-sum) *Consider a zero-sum game* (S_1, S_2, p_1, p_2). *For* $i = 1, 2$ *we have*

$$maxmin_i = -minmax_{-i}$$

and

$$minmax_i = -maxmin_{-i}.$$

Proof. Fix i. For each joint strategy (s_i, s_{-i})

$$p_i(s_i, s_{-i}) = -p_{-i}(s_i, s_{-i}),$$

so

$$\max_{s_i} \min_{s_{-i}} p_i(s_i, s_{-i}) = \max_{s_i} (\min_{s_{-i}} -p_{-i}(s_i, s_{-i})) = -\min_{s_i} \max_{s_{-i}} p_{-i}(s_i, s_{-i}).$$

This proves the first equality. By interchanging i and $-i$ we get the second equality. □

It follows by the Strictly Competitive Games Theorem 18(i) that for zero-sum games a Nash equilibrium exists iff $maxmin_1 = minmax_1$. When this condition holds in a zero-sum game, any pair of security strategies for both players is called a ***saddle point*** of the game and the common value of $maxmin_1$ and $minmax_1$ is called the ***value*** of the game.

Example 15 To illustrate the introduced concepts consider the zero-sum game represented by the following reward matrix:

	L	M	R
T	4	3	5
B	6	2	1

To compute $maxmin_1$ and $minmax_1$, as in Example 14, we extend the matrix with an additional row and column and fill in the minima of the rows and the maxima of the columns:

	L	M	R	f_1
T	4	3	5	3
B	6	2	1	1
F_1	6	3	5	

We see that $maxmin_1 = minmax_1 = 3$. So 3 is the value of this game. Moreover, (T, M) is the only pair of the security strategies of the row and column players, i.e., the only saddle point in this game. □

The above result does not hold for arbitrary strictly competitive games. To see it notice that in any two-player game a multiplication of the payoffs of player i by 2 leads to the doubling of the value of $maxmin_i$ and it does not affect the value of $minmax_{-i}$. Moreover, this multiplication procedure does not affect the property that a game is strictly competitive.

In an arbitrary strategic game with multiple Nash equilibria, for example the Battle of the Sexes game, the players face the following coordination problem. Suppose that each of them chooses a strategy from a Nash equilibrium. Then it can happen that this way they selected a joint strategy that is not a Nash equilibrium. For instance, in the Battle of the Sexes game the players can choose respectively F and B. The following result shows that in a zero-sum game such a coordination problem does not exist.

Theorem 20 (Interchangeability) *Consider a zero-sum game G.*

(i) Suppose that a Nash equilibrium of G exists. Then any joint strategy (s_1^*, s_2^*) such that each s_i^* is a security strategy for player i is a Nash equilibrium of G.

(ii) Suppose that (s_1^*, s_2^*) and (t_1^*, t_2^*) are Nash equilibria of G. Then so are (s_1^*, t_2^*) and (t_1^*, s_2^*).

Proof.
(i) Let (s_1^*, s_2^*) be a pair of security strategies for players 1 and 2. Fix i. By definition

$$\min_{s_i} p_{-i}(s_i, s_{-i}^*) = maxmin_{-i}. \tag{6.3}$$

But

$$\min_{s_i} p_{-i}(s_i, s_{-i}^*) = \min_{s_i} -p_i(s_i, s_{-i}^*) = -\max_{s_i} p_i(s_i, s_{-i}^*)$$

and by the Zero-sum Theorem 19

$$maxmin_{-i} = -minmax_i.$$

So (6.3) implies

$$\max_{s_i} p_i(s_i, s_{-i}^*) = minmax_i. \tag{6.4}$$

We now rely on the Strictly Competitive Games Theorem 18. By item (ii) for $j = 1, 2$ we have $maxmin_j = minmax_j$, so by the proof of item (i) and (6.4) we conclude that (s_i^*, s_{-i}^*) is a Nash equilibrium.

(ii) By (i) and the Strictly Competitive Games Theorem 18(iv). $\square$

The assumption that a Nash equilibrium exists is obviously necessary in item (i) of the above theorem. Indeed, in the finite zero-sum games security strategies always exist, in contrast to the Nash equilibrium.

Finally, recall that throughout this chapter we assumed the existence of various minima and maxima. So the results of this chapter apply only to a specific class of strictly competitive and zero-sum games. This class includes finite games. We shall return to this matter in a later chapter.

Chapter 7

Sealed-bid Auctions

An ***auction*** is a procedure used for selling and buying items by offering them up for bid. Auctions are often used to sell objects that have a variable price (for example oil) or an undetermined price (for example radio frequencies). There are several types of auctions. In its most general form they can involve multiple buyers and multiple sellers with multiple items being offered for sale, possibly in succession. Moreover, some items can be sold in fractions, for example oil.

Here we shall limit our attention to a simple situation in which only one seller exists and offers one object for sale that has to be sold in its entirety (for example a painting). So in this case an auction is a procedure that involves

- one seller who offers an object for sale,
- n bidders, each bidder i having a valuation $v_i \geq 0$ of the object.

The procedure we discuss here involves submission of ***sealed bids***. More precisely, the bidders simultaneously submit their bids in closed envelopes and the object is allocated, in exchange for a payment, to the bidder who submitted the highest bid (the ***winner***). Such an auction is called a ***sealed-bid auction***. To keep things simple we assume that when more than one bidder submitted the highest bid the object is allocated to the highest bidder with the lowest index.

To formulate a sealed-bid auction as a strategic game we consider each bidder as a player. Then we view each bid of player i as his possible strategy. We allow any nonnegative real as a bid.

We assume that the valuations v_i are fixed and publicly known. This is an unrealistic assumption to which we shall return in a later chapter. However, this assumption is necessary, since the valuations are used in the definition of the payoff functions and by assumption the players have common knowledge of the game and hence of each others' payoff functions. When defining the payoff functions we consider two options, each being determined by the underlying payment procedure.

Given a sequence $b := (b_1, \ldots, b_n)$ of reals, we denote the least l such that $b_l = \max_{k \in \{1,\ldots,n\}} b_k$ by $\text{argsmax}\, b$. That is, $\text{argsmax}\, b$ is the smallest index l such that b_l is a largest element in the sequence b. For example, $\text{argsmax}\,(6, 7, 7, 5) = 2$.

7.1 First-price auction

The most commonly used rule in a sealed-bid auction is that the winner i pays to the seller the amount equal to his bid. The resulting mechanism is called the ***first-price auction***.

Assume the winner is bidder i, whose bid is b_i. Since his value for the sold object is v_i, his payoff (profit) is $v_i - b_i$. For the other players the payoff (profit) is 0. Note that the winner's profit can be negative. This happens when he wins the object by ***overbidding***, i.e., submitting a bid higher than his valuation of the object being sold. Such a situation is called the ***winner's curse***.

To summarize, the payoff function p_i of player i in the game associated with the first-price auction is defined as follows, where b is the vector of the submitted bids:

$$p_i(b) := \begin{cases} v_i - b_i & \text{if } i = \text{argsmax}\, b \\ 0 & \text{otherwise} \end{cases}$$

Let us now analyze the resulting game. The following theorem provides a complete characterization of its Nash equilibria.

Theorem 21 (Characterization I) *Consider the game associated with the first-price auction with the players' valuations v. Then b is a Nash equilibrium iff for $i = \text{argsmax}\, b$*

(i) $b_i \leq v_i$

(the winner does not suffer from the winner's curse),

(ii) $\max_{j \neq i} v_j \leq b_i$

(the winner submitted a sufficiently high bid),

(iii) $b_i = \max_{j \neq i} b_j$

(another player submitted the same bid as player i).

These three conditions can be compressed into the single statement

$$\max_{j \neq i} v_j \leq \max_{j \neq i} b_j = b_i \leq v_i,$$

where $i = \text{argsmax} b$. Also note that (i) and (ii) imply that $v_i = \max v$, which means that in every Nash equilibrium a player with the highest valuation is the winner.

Proof.

$(\Rightarrow)$

(i) If $b_i > v_i$, then player's i payoff is negative and it increases to 0 if he submits the bid equal to v_i.

(ii) If $\max_{j \neq i} v_j > b_i$, then player j such that $v_j > b_i$ can win the object by submitting a bid in the open interval (b_i, v_j), say $v_j - \epsilon$. Then his payoff increases from 0 to ϵ.

(iii) If $b_i > \max_{j \neq i} b_j$, then player i can increase his payoff by submitting a bid in the open interval $(\max_{j \neq i} b_j, b_i)$, say $b_i - \epsilon$. Then his payoff increases from $v_i - b_i$ to $v_i - b_i + \epsilon$.

So if any of the conditions $(i) - (iii)$ is violated, then b is not a Nash equilibrium.

$(\Leftarrow)$ Suppose that a vector of bids b satisfies $(i) - (iii)$. Player i is the winner and by (i) his payoff is non-negative. His payoff can increase only if he bids less, but then by (iii) another player (the one who initially submitted the same bid as player i) becomes the winner, while player's i payoff becomes 0.

The payoff of any other player j is 0 and can increase only if he bids more and becomes the winner. But then by (ii), $\max_{j \neq i} v_j < b_j$, so his payoff becomes negative.

So b is a Nash equilibrium. □

As an illustration of the above theorem suppose that the vector of the valuations is $(1, 6, 5, 2)$. Then the vectors of bids $(1, 5, 5, 2)$ and $(1, 5, 2, 5)$ satisfy the above three conditions and are both Nash equilibria. The first vector of bids shows that player 2 can secure the object by bidding the second highest valuation. In the second vector of bids player 4 overbids but his payoff is 0 since he is not the winner.

By the ***truthful bidding*** we mean the vector b of bids, such that for each player i we have $b_i = v_i$, i.e., each player bids his own valuation. Note that by the Characterization Theorem 21 truthful bidding, i.e., v, is a Nash equilibrium iff the two highest valuations coincide.

Further, note that for no player i such that $v_i > 0$ his truthful bidding is a dominant strategy. Indeed, truthful bidding by player i always results in payoff 0. However, if all other players bid 0, then player i can increase his payoff by submitting a lower, positive bid.

Observe that the above analysis does not allow us to conclude that in each Nash equlibrium the winner is the player who wins in the case of truthful bidding. Indeed, suppose that the vector of valuations is $(0, 5, 5, 5)$, so that in the case of truthful bidding by all players player 2 is the winner. Then the vector of bids $(0, 4, 5, 5)$ is a Nash equilibrium with player 3 being the winner.

Finally, notice the following strange consequence of the above theorem: in no Nash equilibrium the last player, n, is a winner. The reason is that we resolved the ties in the favour of a bidder with the lowest index. Indeed, by item (*iii*) in every Nash equilibrium b we have $\text{argsmax}\, b < n$.

7.2 Second-price auction

We consider now an auction with the following payment rule. As before the winner is the bidder who submitted the highest bid (with a tie broken, as before, to the advantage of the bidder with the smallest index), but now he pays to the seller the amount equal to the *second* highest bid. This sealed-bid auction is called the ***second-price auction***. It was proposed by W. Vickrey and is alternatively called ***Vickrey auction***. So in this auction in the absence of ties the winner pays to the seller a lower price than in the first-price auction.

Let us formalize this auction as a game. The payoffs are now defined as follows:

$$p_i(b) := \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } i = \text{argsmax}\, b \\ 0 & \text{otherwise} \end{cases}$$

Note that bidding v_i always yields a non-negative payoff but can now lead to a strictly positive payoff, which happens when v_i is a unique winning bid. However, when the highest two bids coincide the payoffs are still the same as in the first-price auction, since then for $i = \text{argsmax} b$ we have $b_i = \max_{j \neq i} b_j$. Finally, note that the winner's curse still can take place here, namely when $v_i < b_i$ and some other bid is in the open interval (v_i, b_i).

The analysis of the second-price auction as a game leads to different conclusions that for the first-price auction. The following theorem provides a complete characterization of the Nash equilibria of the corresponding game.

Theorem 22 (Characterization II) *Consider the game associated with the second-price auction with the players' valuations v. Then b is a Nash equilibrium iff for $i = argsmax\, b$*

(i) $\max_{j \neq i} v_j \leq b_i$

(the winner submitted a sufficiently high bid),

(ii) $\max_{j \neq i} b_j \leq v_i$

(the winner's valuation is sufficiently high).

Proof.

$(\Rightarrow)$

(i) If $\max_{j \neq i} v_j > b_i$, then player j such that $v_j > b_i$ can win the object by submitting a bid in the open interval (b_i, v_j). Then his payoff increases from 0 to $v_j - b_i$.

(ii) If $\max_{j \neq i} b_j > v_i$, then player's i payoff is negative, namely $v_i - \max_{j \neq i} b_j$, and can increase to 0 if player i submits a losing bid.

So if condition (i) or (ii) is violated, then b is not a Nash equilibrium.

$(\Leftarrow)$ Suppose that a vector of bids b satisfies (i) and (ii). Player i is the winner and by (ii) his payoff is non-negative. By submitting another bid he either remains a winner, with the same payoff, or becomes a loser with the payoff 0.

The payoff of any other player j is 0 and can increase only if he bids more and becomes the winner. But then his payoff becomes $v_j - b_i$, so by (i) becomes negative.

So b is a Nash equilibrium. $\square$

This characterization result shows that several Nash equilibria exist. We now exhibit three specific ones that are of particular interest. In each case it is straightforward to check that conditions (i) and (ii) of the above theorem hold.

Truthful bidding

Recall that in the case of the first-price auction truthful bidding is a Nash equilibrium iff for the considered sequence of valuations the auction coincides with the second-price auction. Now truthful bidding, so v, is always a Nash equilibrium. Below we prove another property of truthful bidding in second-price auction.

Wolf and sheep Nash equilibrium

Suppose that $i = \text{argsmax}\, v$, i.e., player i is the winner in the case of truthful bidding. Consider the strategy profile in which player i bids v_i and everybody else bids 0. This Nash equilibrium is called ***wolf and sheep***, where player i plays the role of a wolf by bidding aggressively and scaring the sheep being the other players who submit their minimal bids.

Yet another Nash equilibrium

Finally, we exhibit a Nash equilibrium in which the player with the uniquely highest valuation is not a winner. This is in contrast with what we observed in the case of the first-price auction. Suppose that the two highest bids are v_j and v_i, where $i < j$ and $v_j > v_i > 0$. Then the strategy profile in which player i bids v_j, player j bids v_i and everybody else bids 0 is a Nash equilibrium.

In both the first-price and the second-price auctions overbidding, i.e., submitting a bid above one's valuation of the object looks risky and therefore not credible. Note that the bids that do not exceed one's valuation are exactly the security strategies.

So when we add the following additional condition to each characterization theorem:

- for all $j \in \{1, \ldots, n\}$, $b_j \leq v_j$,

we characterize in each case Nash equilibria in the security strategies.

7.3 Incentive compatibility

So far we discussed two examples of sealed-bid auctions. A general form of such an auction is determined by fixing for each bidder i the payment procedure pay_i which given a sequence b of bids such that bidder i is the winner yields his payment.

In the resulting game, that we denote by $G_{pay,v}$, the payoff function is defined by

$$p_i(b) := \begin{cases} v_i - pay_i(b) & \text{if } i = \text{argsmax}\, b \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, bidding 0 means that the bidder is not interested in the object. So if all players bid 0 then none of them is interested in the object. According to our definition the object is then allocated to the first bidder. We assume that then his payment is 0. That is, we stipulate that $pay_1(0, \ldots, 0) = 0$.

When designing a sealed-bid auction it is natural to try to induce the bidders to bid their valuations. This leads to the following notion.

We call a sealed-bid auction with the payment procedures $pay_1, \ldots, pay_n$ ***incentive compatible*** if for all sequences v of players' valuations for each bidder i his valuation v_i is a dominant strategy in the corresponding game $G_{pay,v}$.

While dominance of a strategy does not guarantee that a player will choose it, it ensures that deviating from it is not profitable. So dominance of each valuation v_i can be viewed as a statement that in the considered auction lying does not pay off.

We now show that the condition of incentive compatibility fully characterizes the corresponding auction. More precisely, the following result holds.

Theorem 23 (Second-price auction) *A sealed-bid auction is incentive compatible iff it is the second-price auction.*

Proof. Fix a sequence of the payment procedures $pay_1, \ldots, pay_n$ that determines the considered sealed-bid auction.

$(\Rightarrow)$ Choose an arbitrary sequence of bids that for the clarity of the argument we denote by (v_i, b_{-i}). Suppose that $i = \text{argsmax}\,(v_i, b_{-i})$. We establish the following four claims.

Claim 1. $pay_i(v_i, b_{-i}) \leq v_i$.
Proof. Suppose by contradiction that $pay_i(v_i, b_{-i}) > v_i$. Then in the corresponding game $G_{pay,v}$ we have $p_i(v_i, b_{-i}) < 0$. On the other hand $p_i(0, b_{-i}) \geq 0$. Indeed, if $i \neq \text{argsmax}\,(0, b_{-i})$, then $p_i(0, b_{-i}) = 0$. Otherwise all bids in b_{-i} are 0 and $i = 1$, and hence $p_i(0, b_{-i}) = v_i$, since by assumption $pay_1(0, \ldots, 0) = 0$.

This contradicts the assumption that v_i is a dominant strategy in the corresponding game $G_{pay,v}$.

Claim 2. For all $b_i \in (\max_{j \neq i} b_j, v_i)$ we have $pay_i(v_i, b_{-i}) \leq pay_i(b_i, b_{-i})$.
Proof. Suppose by contradiction that for some $b_i \in (\max_{j \neq i} b_j, v_i)$ we have $pay_i(v_i, b_{-i}) > pay_i(b_i, b_{-i})$. Then $i = \text{argsmax}\,(b_i, b_{-i})$ so

$$p_i(v_i, b_{-i}) = v_i - pay_i(v_i, b_{-i}) < v_i - pay_i(b_i, b_{-i}) = p_i(b_i, b_{-i}).$$

This contradicts the assumption that v_i is a dominant strategy in the corresponding game $G_{pay,v}$.

Claim 3. $pay_i(v_i, b_{-i}) \leq \max_{j \neq i} b_j$.
Proof. Suppose by contradiction that $pay_i(v_i, b_{-i}) > \max_{j \neq i} b_j$. Take some $v'_i \in (\max_{j \neq i} b_j, pay_i(v_i, b_{-i}))$. By Claim 1 $v'_i < v_i$, so by Claim 2 $pay_i(v_i, b_{-i}) \leq pay_i(v'_i, b_{-i})$. Further, by Claim 1 for the sequence (v'_i, v_{-i}) of valuations we have $pay_i(v'_i, b_{-i}) \leq v'_i$.

So $pay_i(v_i, b_{-i}) \leq v'_i$, which contradicts the choice of v'_i.

Claim 4. $pay_i(v_i, b_{-i}) \geq \max_{j \neq i} b_j$.
Proof. Suppose by contradiction that $pay_i(v_i, b_{-i}) < \max_{j \neq i} b_j$. Take an arbitrary $v'_i \in (pay_i(v_i, b_{-i}), \max_{j \neq i} b_j)$. Then $p_i(v'_i, b_{-i}) = 0$, while

$$p_i(v_i, b_{-i}) = v_i - pay_i(v_i, b_{-i}) > v_i - \max_{j \neq i} b_j \geq 0.$$

This contradicts the assumption that v'_i is a dominant strategy in the corresponding game $G_{pay,(v'_i, v_{-i})}$.

So we proved that for $i = \text{argsmax}\,(v_i, b_{-i})$ we have $pay_i(v_i, b_{-i}) = \max_{j \neq i} b_j$, which shows that the considered sealed-bid auction is second price.

($\Leftarrow$) We actually prove a stronger claim, namely that all sequences of valuations v each v_i is a weakly dominant strategy for player i.

To this end take a vector b of bids. By definition $p_i(b_i, b_{-i}) = 0$ or $p_i(b_i, b_{-i}) = v_i - \max_{j \neq i} b_j \leq p_i(v_i, b_{-i})$. But $0 \leq p_i(v_i, b_{-i})$, so

$$p_i(b_i, b_{-i}) \leq p_i(v_i, b_{-i}).$$

Consider now a bid $b_i \neq v_i$. If $b_i < v_i$, then take b_{-i} such that each element of it lies in the open interval (b_i, v_i). Then b_i is a losing bid and v_i is a winning bid and

$$p_i(b_i, b_{-i}) = 0 < v_i - \max_{j \neq i} b_j = p_i(v_i, b_{-i}).$$

If $b_i > v_i$, then take b_{-i} such that each element of it lies in the open interval (v_i, b_i). Then b_i is a winning bid and v_i is a losing bid and

$$p_i(b_i, b_{-i}) = v_i - \max_{j \neq i} b_j < 0 = p_i(v_i, b_{-i}).$$

So we proved that each strategy $b_i \neq v_i$ is weakly dominated by v_i, i.e., that v_i is a weakly dominant strategy. As an aside, recall that each weakly dominant strategy is unique, so we characterized bidding one's valuation in the second-price auction in game theoretic terms. $\Box$

Exercise 9 Prove that the game associated with the first-price auction with the players' valuations v has no Nash equilibrium iff v_n is the unique highest valuation. $\Box$

Chapter 8

Repeated Games

In the games considered so far the players took just a single decision: a strategy they selected. In this chapter we consider a natural idea of playing a given strategic game repeatedly. We assume that the outcome of each round is known to all players before the next round of the game takes place.

8.1 Finitely repeated games

In the first approach we shall assume that the same game is played a fixed number of times. The final payoff to each player is simply the sum of the payoffs obtained in each round.

Suppose for instance that we play the Prisoner's Dilemma game, so

	C	D
C	$2,2$	$0,3$
D	$3,0$	$1,1$

twice. It seems then that the outcome is the following game in which we simply add up the payoffs from the first and second round:

	CC	CD	DC	DD
CC	$4,4$	$2,5$	$2,5$	$0,6$
CD	$5,2$	$3,3$	$3,3$	$1,4$
DC	$5,2$	$3,3$	$3,3$	$1,4$
DD	$6,0$	$4,1$	$4,1$	$2,2$

However, this representation is incorrect since it erronously assumes that the decisions taken by the players in the first round have no influence on their decisions taken in the second round. For instance, the option that the first player chooses C in the second round if and only iff the second player chose C in the first round is not listed. In fact, the set of strategies available to each player is much larger.

In the first round each player has two strategies. However, in the second round each player's strategy is a function $f : \{C, D\} \times \{C, D\} \rightarrow \{C, D\}$. So in the second round each player has $2^4 = 16$ strategies and consequently in the repeated game each player has $2 \times 16 = 32$ strategies. Each such strategy has two components, one of each round. It is clear how to compute the payoffs for so defined strategies. For instance, if the first player chooses in the first round C and in the second round the function

$$f_1(s) := \begin{cases} C & \text{if } s = (C, C) \\ D & \text{if } s = (C, D) \\ C & \text{if } s = (D, C) \\ D & \text{if } s = (D, D) \end{cases}$$

and the second player chooses in the first round D and in the second round the function

$$f_2(s) := \begin{cases} C & \text{if } s = (C, C) \\ D & \text{if } s = (C, D) \\ D & \text{if } s = (D, C) \\ C & \text{if } s = (D, D) \end{cases}$$

then the corresponding payoffs are:

- in the first round: $(0, 3)$ (corresponding to the joint strategy (C, D)),
- in the second round: $(1, 1)$ (corresponding to the joint strategy (D, D)).

So the overall payoffs are: $(1, 4)$, which corresponds to the joint strategy (CD, DD) in the above bimatrix.

Let us consider now the general setup. The strategic game that is repeatedly played is called the ***stage game***. Given a stage game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ the ***repeated game with k rounds*** (in short: a repeated game), where $k \geq 1$, is defined by first introducing the set of ***histories***. This set $\mathcal{H}$ is defined inductively as follows, where ε denotes the empty sequence, $t \geq 1$, and, as usual, $S = S_1 \times \ldots \times S_n$:

$$\begin{aligned}
&\mathcal{H}^0 := \{\varepsilon\}, \\
&\mathcal{H}^1 := S, \\
&\mathcal{H}^{t+1} := \mathcal{H}^t \times S, \\
&\mathcal{H} := \bigcup_{t=0}^{k-1} \mathcal{H}^t.
\end{aligned}$$

So $h \in \mathcal{H}^0$ iff $h = \varepsilon$ and for $t \in \{1, \ldots, k-1\}$, $h \in \mathcal{H}^t$ iff $h \in S^t$. That is, a history is a (possibly empty) sequence of joint strategies of the stage game of length at most $k-1$.

Then a ***strategy*** for player i in the repeated game is a function $\sigma_i : \mathcal{H} \to S_i$. In particular $\sigma_i(\varepsilon)$ is a strategy in the stage game chosen in the first round.

We denote the set of strategies of player i in the repeated game by Σ_i and the set of joint strategies in the repeated game by Σ.

The ***outcome*** of the repeated game corresponding to a joint strategy $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma$ of the players is the history that consists of k joint strategies selected in the consecutive stages of the underlying stage game. This history $(o^1(\sigma), \ldots, o^k(\sigma)) \in \mathcal{H}^k$ is defined as follows:

$$\begin{aligned}
&o^1(\sigma) := (\sigma_1(\varepsilon), \ldots, \sigma_n(\varepsilon)), \\
&o^2(\sigma) := (\sigma_1(o^1(\sigma)), \ldots, \sigma_n(o^1(\sigma))), \\
&\ldots \\
&o^k(\sigma) := (\sigma_1(o^1(\sigma), \ldots, o^{k-1}(\sigma)), \ldots, \sigma_n(o^1(\sigma), \ldots, o^{k-1}(\sigma)).
\end{aligned}$$

In particular $o^k(\sigma)$ is obtained by applying each of the strategies $\sigma_1, \ldots, \sigma_n$ to the already defined history $(o^1(\sigma), \ldots, o^{k-1}(\sigma)) \in \mathcal{H}^{k-1}$.

Finally, the ***payoff function*** P_i of player i in the repeated game is defined as

$$P_i(\sigma) := \sum_{t=1}^{k} p_i(o^t(\sigma)).$$

So the payoff for each player is the sum of the payoffs he received in each round.

Now that we defined formally a repeated game let us return to the Prisoner's Dilemma game and assume that it is played k rounds. We can now define the following natural strategies:[1]

[1] These definitions are incomplete in the sense that the strategies are not defined for all

- ***cooperate***: select at every stage C,
- ***defect***: select at every stage D,
- ***tit for tat***: first select C, then repeatly select the last strategy played by the opponent,
- ***grim*** (or ***trigger***): select C as long as the opponent selects C; if he selects D select D from now on.

For example, it does not matter if one chooses tit for tat or grim strategy against a grim strategy: in both cases each player repeatedly selects C. However, if one selects C in the odd rounds and D in the even rounds, then against the tit for tat strategy the following sequence of stage strategies results:

- for player 1: $C, D, C, D, C, \ldots$,
- for player 2: $C, C, D, C, D, \ldots$

while against the grim strategy we obtain:

- for player 1: $C, D, C, D, C, \ldots$,
- for player 2: $C, C, D, D, D, \ldots$

Using the concept of strictly dominant strategies we could predict that the outcome of the Prisoner's dilemma game is (D, D). A natural question arises whether we can also predict the outcome in the repeated version of this game. To do this we first extend the relevant notions to the repeated games.

Given a stage game G we denote the repeated game with k rounds by $G(k)$. After the obvious identification of $\sigma_i : \mathcal{H}^0 \rightarrow S_i$ with $\sigma_i(\varepsilon)$ we can identify $G(1)$ with G.

In general we can view $G(k)$ as a strategic game $(\Sigma_1, \ldots, \Sigma_n, P_1, \ldots, P_n)$, where the strategy sets Σ_i and the payoff functions P_i are defined above. This allows us to apply the basic notions, for example that of Nash equilibrium, to the repeated game.

As a first result we establish the following.

histories. However, the specified parts completely determine the outcomes that can arise against any strategy of the opponent.

Theorem 24 *Consider a stage game G and $k \geq 1$.*

(i) If s is a Nash equilibrium of G, then the joint strategy σ, where for all $i \in \{1, \ldots, n\}$ and $h \in \mathcal{H}$

$$\sigma_i(h) := s_i,$$

is a Nash equilibrium of $G(k)$.

(ii) If s is a unique Nash equilibrium of G, then for each Nash equilibrium of $G(k)$ the outcome corresponding to it consists of s repeated k times.

Proof.

(i) The outcome corresponding to σ consists of s repeated k times. That is, in each round of $G(k)$ the Nash equilibrium is selected and the payoff to each player i is $p_i(s)$, where $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$.

Suppose that σ is not a Nash equilibrium in $G(k)$. Then for some player i a strategy τ_i yields a higher payoff than σ_i when used against σ_{-i}. So in some round of $G(k)$ player i receives a strictly larger payoff than $p_i(s)$. But in this (and every other) round every other player j selects s_j. So the strategy of player i selected in this round yields a strictly higher payoff against s_{-i} than s_i, which is a contradiction.

(ii) We proceed by induction on k. Since we identified $G(1)$ with G, the claim holds for $k = 1$. Suppose it holds for $k \geq 1$.

Take a Nash equilibrium σ in $G(k+1)$. Consider the last joint strategy of the considered outcome. It constitutes the Nash equilibrium of the stage game. Indeed, otherwise some player i did not select a best response in the last round and thus can obtain a strictly higher payoff in G by switching in the last round of $G(k+1)$ to another strategy s'_i in G. The corresponding modification of σ_i according to which in the last round s'_i is selected yields against σ_{-i} a strictly higher payoff than σ_i. This contradicts the assumption that σ is a Nash equilibrium in $G(k+1)$.

Now redistribute for each player his payoff in the last round evenly over the previous k rounds, by modifying appropriately the payoff functions, and subsequently remove this last round. The resulting game is a repeated game $G'(k)$ such that s is a unique Nash equilibrium of G', so we can apply to it the induction hypothesis. Moreover, by the above observation each Nash equilibrium of $G(k+1)$ consists of a Nash equilibrium of $G'(k)$ augmented with s selected in the last round (i.e., in each Nash equilibrium of $G(k+1)$ each player i selects s_i in the last round). So the claim holds for $k+1$.

The claim now holds by induction. □

The definition of a strategy in a repeated game determines player's choice for *each* history, in particular for histories that cannot be outcomes of the repeated game. As a result the joint strategy from item (i) is not a unique Nash equilibrium of $G(k)$ when players have two or more strategies in the stage game.

As an example consider the Prisoner's Dilemma game played twice. Then the pair of defect strategies is a Nash equilibrium. Moreover, the pair of strategies according to which one selects D in the first round and C in the second round iff the first round equals (C, C) is also a Nash equilibrium. These two pairs differ though they yield the same outcome.

Note, further that if a player has a strictly dominant strategy in the stage game then he does not necessarily have a strictly dominant strategy in the repeated game. In particular, choosing in each round the strictly dominant strategy in the stage game does not need to yield a maximal payoff in the repeated game.

Example 16 Take the Prisoner's Dilemma game played twice.

Consider first a best response against the tit for tat strategy. In it C is selected in the first round and D in the second round. In contrast, in each best response against the cooperate strategy in both rounds D is selected. So for each player no single best response strategy exists, that is, no player has a strictly dominant strategy.

In contrast, in the stage game strategy D is strictly dominant for both players. Note also that in our first, incorrect, representation of the Prisoner's Dilemma game played twice strategy DD is strictly dominant for both players, as well. □

In the one shot version of the Prisoner's Dilemma game we could predict that both players will select the defect (D) strategy on the basis that it is a strictly dominant strategy. The above theorem shows that when Prisoner's Dilemma game is played repeatedly cooperation still won't occur. However, this prediction is weaker, in the sense that selecting D repeatedly is not anymore a strictly dominant strategy. We can only conclude that in any Nash equilibrium in every round each player selects D.

The above theorem actually shows more: when the stage game has exactly one Nash equilibrium, then in each Nash equilibrium of the repeated game the players select their equilibrium strategies. So in each round their payoff is simply their payoff in the Nash equilibrium of the stage game.

However, when the stage game has more than one Nash equilibrium the situation changes. In particular, players can achieve in a Nash equilibrium of the repeated game a higher average payoff than the one achieved in any Nash equilibrium of the stage game.

Example 17 Consider the following stage game:

	A	B	C
A	$5,5$	$0,0$	$12,0$
B	$0,0$	$2,2$	$0,0$
C	$0,12$	$0,0$	$10,10$

This game has two Nash equilibria (A, A) and (B, B). So when the game is played once the highest payoff in a Nash equilibrium is 5 for each player. However, when the game is played twice a Nash equilibrium exists with a higher average payoff. Namely, consider the following strategy for each player:

- select C in the first round,
- if the other player selected C in the first round, select A in the second round and otherwise select M.

If each player selects this strategy, they both select in the first round C and A in the second round. This yields payoff 15 for each player.

We now prove that this pair of strategies forms a Nash equlibrium. The only way a player, say the first one, can receive a larger payoff than 15 is by selecting A in the first round. But then the second player selects B in the second round. So in the first round the first player receives the payoff 12 but in the second round he receives the payoff of at most 2. Consequently by switching to another strategy the first player can secure at best payoff 14. □

The above example shows that playing a given game repeatedly can lead to some form of coordination that can be beneficial to all players. This coordination is possible because crucially the choices made by the players in the previous rounds are commonly known.

Exercise 10 Compute the strictly and weakly dominated strategies in the Prisoner's Dilemma game played twice. □

8.2 Infinitely repeated games

In this section we consider infinitely repeated games. To define them we need to modify appropriately the approach of the previous section.

First, to ensure that the payoffs are well defined we assume that in the underlying stage game the payoff functions are bounded (from above and below). Then we redefine the set of ***histories*** by putting

$$\mathcal{H} := \bigcup_{t=0}^{\infty} \mathcal{H}^t,$$

where each $\mathcal{H}^t$ is defined as before.

The notion of a ***strategy*** of a player remains the same: it is a function from the set of all histories to the set of his strategies in the stage game. An ***outcome*** corresponding to a joint strategy σ is now the infinite set of joint strategies of the stage game $o^1(\sigma), o^2(\sigma), \ldots$ where each $o^t(\sigma)$ is defined as before.

Finally, to define the ***payoff function*** we first introduce a ***discount***, which is a number $\delta \in (0, 1)$. Then we put

$$P_i(\sigma) := (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} p_i(o^t(\sigma)).$$

This definition requires some explanations. First note that this payoff function is well-defined and always yields a finite value. Indeed, the original payoff functions are assumed to be bounded and $\delta \in (0, 1)$, so the sequence $(\sum_{t=1}^{t} \delta^{t-1} p_i(o^t(\sigma)))_{t=1,2,\ldots}$ converges.

Note that the payoff in each round t is discounted by δ^{t-1}, which can be viewed as the accumulated depreciation. So discounted payoffs in each round are summed up and subsequently multiplied by the factor $1 - \delta$. Note that

$$\sum_{t=1}^{\infty} \delta^{t-1} = 1 + \delta \sum_{t=1}^{\infty} \delta^{t-1},$$

hence

$$\sum_{t=1}^{\infty} \delta^{t-1} = \frac{1}{1 - \delta}.$$

So if in each round the players select the same joint strategy s, then their respective payoffs in the stage game and the repeated game coincide. This explains the adjustment factor $1-\delta$ in the definition of the payoff functions. Further, since the payoffs in the stage game are bounded, the payoffs in the repeated game are finite.

Given a stage game G and a discount δ we denote the infinitely repeated game defined above by $G(\delta)$.

We observed in the previous section that in each Nash equilibrium of the finitely repeated Prisoner's Dilemma game the players select in each round the defect (D) strategy. So finite repetition does not allow us to induce cooperation, i.e., the selection of the C strategy. We now show that in the infinitely repeated game the situation dramatically changes. Namely, the following holds.

Theorem 25 (Prisoner's Dilemma) *Take as G the Prisoner's Dilemma game. Then for all $\delta \in (\frac{1}{2}, 1)$ the pair of trigger strategies forms a Nash equilibrium of $G(\delta)$.*

Note that the outcome corresponding to the pair of trigger strategies consists of the infinite sequence of (C, C), that is, in the claimed Nash equilibrium of $G(\delta)$ both players repeatedly select C, i.e., always cooperate.

Proof. Suppose that, say, the first player deviates from his trigger strategy while the other player remains at his trigger strategy. Let t be the first stage in which the first player selects D. Consider now his payoffs in the consecutive rounds of the stage game:

- in the rounds $1, \ldots, t-1$ they equal 2,
- in the round t it equals 3,
- in the rounds $t+1, \ldots,$ they equal at most 1.

So the payoff in the repeated game is bounded from above by

$$\begin{array}{rl} & (1-\delta)(2\sum_{j=1}^{t-1} \delta^{j-1} + 3\delta^{t-1} + \sum_{j=t+1}^{\infty} \delta^{j-1}) \\ = & (1-\delta)(2\frac{1-\delta^{t-1}}{1-\delta} + 3\delta^{t-1} + \frac{\delta^t}{1-\delta}) \\ = & 2(1-\delta^{t-1}) + 3\delta^{t-1}(1-\delta) + \delta^t \\ = & 2 + \delta^{t-1} - 2\delta^t. \end{array}$$

Since $\delta > 0$, we have

$$\delta^{t-1} - 2\delta^t < 0 \text{ iff } 1 - 2\delta < 0 \text{ iff } \delta > \tfrac{1}{2}.$$

So when the first player deviates from his trigger strategy and $\delta > \frac{1}{2}$, his payoff in the repeated game is less than 2. In contrast, when he remains at the trigger strategy, his payoff is 2.

This concludes the proof. $\square$

This theorem shows that cooperation can be achieved by repeated interaction, so it seems to carry a positive message. However, repeated selection of the defect strategy D by both players still remains a Nash equilibrium and there is an obvious coordination problem between these two Nash equilibria.

Moreover, the above result is a special case of a much more general theorem. To formulate it we shall use the $minmax_i$ value introduced in Section 5.2 that, given a game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ was defined by

$$minmax_i := \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} p_i(s_i, s_{-i}).$$

The following result is called Folk theorem since some version of it has been known before it was recorded in a journal paper. From now on we abbreviate $(p_1(s), \ldots, p_n(s))$ to $p(s)$ and similarly with the P_i payoff functions.

Theorem 26 (Folk Theorem) *Consider a stage game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ with the bounded payoff functions.*

Take some $s' \in S$ and suppose $r := p(s')$ is such that for $i \in \{1, \ldots, n\}$ we have $r_i > minmax_i$. Then $\delta_0 \in (0, 1)$ exists such that for all $\delta \in (\delta_0, 1)$ the repeated game $G(\delta)$ has a Nash equilibrium σ with $P(\sigma) = r$.

Note that this theorem is indeed a generalization of the Prisoner's Dilemma Theorem 25 since for the Prisoner's Dilemma game we have $minmax_1 = minmax_2 = 1$, while for the joint strategy (C, C) the payoff to each player is 2. Now, the only way to achieve this payoff for both players in the repeated game is by repeatedly selecting C.

Proof. The argument is analogous to the one we used in the proof of the Prisoner's Dilemma Theorem 25. Let the strategy σ_i consist of selecting in each round s'_i. Note that $P(\sigma) = r$.

We first define an analogue of the trigger strategy. Let s^*_{-i} be such that $\max_{s_i} p_i(s_i, s^*_{-i}) = minmax_i$. That is, s^*_{-i} is the joint strategy of the opponents of player i that when selected by them results in a minimum possible payoff to player i. The idea behind the strategies defined below is that the opponents of the deviating player jointly switch forever to s^*_{-i} to 'inflict' on player i the maximum 'penalty'.

Recall that a history h is a finite sequence of joint strategies in the stage game. Below a *deviation in* h refers to the fact that a specific player i did not select s'_i in a joint strategy from h.

Given $h \in \mathcal{H}$ and $j \in \{1, \ldots, n\}$ we put

$$\sigma_j(h) := \begin{cases} s'_j & \text{if no player } i \neq j \text{ deviated in } h \text{ from } s'_i \text{ unilaterally} \\ s^*_j & \text{otherwise, where } i \text{ is the first player who deviated in } h \text{ from} \\ & s'_i \text{ unilaterally} \end{cases}$$

We now claim that σ is a Nash equilibrium for appropriate δs. Suppose that some player i deviates from his strategy σ_i while the other players remain at σ_{-i}. Let t be the first stage in which player i selects a strategy s''_i different from s'_i. Consider now his payoffs in the consecutive rounds of the stage game:

- in the rounds $1, \ldots, t-1$ they equal r_i,
- in the round t it equals $p_i(s''_i, s'_{-i})$,
- in the rounds $t+1, \ldots,$ they equal at most $minmax_i$.

Let $r^*_i > p_i(s)$ for all $s \in S$. The payoff of player i in the repeated game $G(\delta)$ is bounded from above by

$$\begin{array}{rl} & (1-\delta)(r_i \sum_{j=1}^{t-1} \delta^{j-1} + r^*_i \delta^{t-1} + minmax_i \sum_{j=t+1}^{\infty} \delta^{j-1}) \\ = & (1-\delta)(r_i \frac{1-\delta^{t-1}}{1-\delta} + r^*_i \delta^{t-1} + minmax_i \frac{\delta^t}{1-\delta}) \\ = & r_i - \delta^{t-1} r_i + (1-\delta)\delta^{t-1} r^*_i + \delta^t minmax_i \\ = & r_i + \delta^{t-1}(-r_i + (1-\delta) r^*_i + \delta\, minmax_i). \end{array}$$

Since $\delta > 0$ and $r^*_i \geq r_i > minmax_i$, we have

$$\begin{array}{rl} & \delta^{t-1}(-r_i + (1-\delta) r^*_i + \delta\, minmax_i) < 0 \\ \text{iff} & r^*_i - r_i - \delta(r^*_i - minmax_i) < 0 \\ \text{iff} & \frac{r^*_i - r_i}{r^*_i - minmax_i} < \delta. \end{array}$$

But $r^*_i > r_i > minmax_i$ implies that $\delta_0 := \frac{r^*_i - r_i}{r^*_i - minmax_i} \in (0, 1)$. So when $\delta > \delta_0$ and player i selects in some round a strategy different than s'_i, while every other player j keeps selecting s'_j, player's i payoff in the repeated game is less than r_i. In contrast, when he remains selecting s'_i his payoff is r_i.

So σ is indeed a Nash equilibrium. □

The above result can be strengthened to a much larger set of payoffs. Recall that a set of points $A \subseteq \mathbb{R}^n$ is called ***convex*** if for any $\mathbf{x}, \mathbf{y} \in A$ and $\alpha \in [0, 1]$ we have $\alpha\mathbf{x} + (1 - \alpha)\mathbf{y} \in A$. Given a subset $A \subseteq \mathbb{R}^k$ denote the smallest convex set that contains A by $conv(A)$.

Then the above theorem holds not only for $r \in \{p(s) \mid s \in S\}$, but also for all $r \in conv(\{p(s) \mid s \in S\})$. In the case of the Prisoner's Dilemma game G we get that for *any*

$$r \in conv(\{(2, 2), (3, 0), (0, 3), (1, 1)\}) \cap \{r' \mid r'_1 > 1, r'_2 > 1\}$$

there is $\delta_0 \in (0, 1)$ such that for all $\delta \in (\delta_0, 1)$ the repeated game $G(\delta)$ has a Nash equilibrium σ with $P(\sigma) = r$. In other words, cooperation can be achieved in a Nash equilibrium, but equally well many other outcomes.

Such results belong to a class of similar theorems collectively called Folks theorems. The considered variations allow for different sets of payoffs achievable in an equilibrium, different ways of computing the payoff, different forms of equilibria, and different types of repeated games.

Chapter 9

Mixed Extensions

We now study a special case of infinite strategic games that are obtained in a canonic way from the finite games, by allowing mixed strategies. Below $[0,1]$ stands for the real interval $\{r \in \mathbb{R} \mid 0 \leq r \leq 1\}$. By a ***probability distribution*** over a finite non-empty set A we mean a function

$$\pi : A \rightarrow [0,1]$$

such that $\sum_{a \in A} \pi(a) = 1$. We denote the set of probability distributions over A by ΔA.

9.1 Mixed strategies

Consider now a finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$. By a ***mixed strategy*** of player i in G we mean a probability distribution over S_i. So ΔS_i is the set of mixed strategies available to player i. In what follows, we denote a mixed strategy of player i by m_i and a joint mixed strategy of the players by m.

Given a mixed strategy m_i of player i we define

$$support(m_i) := \{a \in S_i \mid m_i(a) > 0\}$$

and call this set the ***support*** of m_i. In specific examples we write a mixed strategy m_i as the sum $\sum_{a \in A} m_i(a) \cdot a$, where A is the support of m_i.

Note that in contrast to S_i the set ΔS_i is infinite. When referring to the mixed strategies, as in the previous chapters, we use the '$_{-i}$' notation. So for $m \in \Delta S_1 \times \ldots \times \Delta S_n$ we have $m_{-i} = (m_j)_{j \neq i}$, etc.

We can identify each strategy $s_i \in S_i$ with the mixed strategy that puts 'all the weight' on the strategy s_i. In this context s_i will be called a ***pure strategy***. Consequently we can view S_i as a subset of ΔS_i and S_{-i} as a subset of $\times_{j \neq i} \Delta S_j$.

By a ***mixed extension*** of $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ we mean the strategic game

$$(\Delta S_1, \ldots, \Delta S_n, p_1, \ldots, p_n),$$

where each function p_i is extended in a canonic way from $S := S_1 \times \ldots \times S_n$ to $M := \Delta S_1 \times \ldots \times \Delta S_n$ by first viewing each joint mixed strategy $m = (m_1, \ldots, m_n) \in M$ as a probability distribution over S, by putting for $s \in S$

$$m(s) := m_1(s_1) \cdot \ldots \cdot m_n(s_n),$$

and then by putting

$$p_i(m) := \sum_{s \in S} m(s) \cdot p_i(s).$$

Example 18 Reconsider the Battle of the Sexes game from Chapter 1. Suppose that player 1 (man) chooses the mixed strategy $\frac{1}{2}F + \frac{1}{2}B$, while player 2 (woman) chooses the mixed strategy $\frac{1}{4}F + \frac{3}{4}B$. This pair m of the mixed strategies determines a probability distribution over the set of joint strategies, that we list to the left of the bimatrix of the game:

	F	B
F	$\frac{1}{8}$	$\frac{3}{8}$
B	$\frac{1}{8}$	$\frac{3}{8}$

	F	B
F	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

To compute the payoff of player 1 for this mixed strategy m we multiply each of his payoffs for a joint strategy by its probability and sum it up:

$$p_1(m) = \frac{1}{8}2 + \frac{3}{8}0 + \frac{1}{8}0 + \frac{3}{8}1 = \frac{5}{8}.$$

Analogously

$$p_2(m) = \frac{1}{8}1 + \frac{3}{8}0 + \frac{1}{8}0 + \frac{3}{8}2 = \frac{7}{8}.$$

□

This example suggests the computation of the payoffs in two-player games using matrix multiplication. First, we view each bimatrix of such a game as a pair of matrices $(\mathbf{A}, \mathbf{B})$. The first matrix represents the payoffs to player 1 and the second one to player 1. Assume now that player 1 has k strategies and player 2 has ℓ strategies. Then both $\mathbf{A}$ and $\mathbf{B}$ are $k \times \ell$ matrices. Further, each mixed strategy of player 1 can be viewed as a row vector $\mathbf{p}$ of length k (i.e., a $1 \times k$ matrix) and each mixed strategy of player 2 as a row vector $\mathbf{q}$ of length ℓ (i.e., a $1 \times \ell$ matrix). Since $\mathbf{p}$ and $\mathbf{q}$ represent mixed strategies, we have $\mathbf{p} \in \Delta^{k-1}$ and $\mathbf{q} \in \Delta^{\ell-1}$, where for all $m \geq 0$

$$\Delta^{m-1} := \{(x_1, \ldots, x_m) \mid \sum_{i=1}^{m} x_i = 1 \text{ and } \forall i \in \{1, \ldots, m\}\ x_i \geq 0\}.$$

Δ^{m-1} is called the $(m-1)$-***dimensional unit simplex***.

In the case of our example we have

$$\mathbf{p} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix}, \mathbf{q} = \begin{pmatrix} \frac{1}{4} & \frac{3}{4} \end{pmatrix}, \mathbf{A} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}, \mathbf{B} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

Now, the payoff functions can be defined as follows:

$$p_1(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^{k} \sum_{j=1}^{\ell} \mathbf{p}_i \mathbf{q}_j \mathbf{A}_{ij} = \mathbf{p}\mathbf{A}\mathbf{q}^T$$

and

$$p_2(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^{k} \sum_{j=1}^{\ell} \mathbf{p}_i \mathbf{q}_j \mathbf{B}_{ij} = \mathbf{p}\mathbf{B}\mathbf{q}^T.$$

9.2 Nash equilibria in mixed strategies

In the context of a mixed extension we talk about a ***pure Nash equlibrium***, when each of the constituent strategies is pure, and refer to an arbitrary Nash equilibrium of the mixed extension as a ***Nash equilibrium in mixed strategies*** of the initial finite game. In what follows, when we use the letter m we implicitly refer to the latter Nash equilibrium.

Below we shall need the following notion. Given a probability distribution π over a finite non-empty multiset[1] A of reals, we call

$$\sum_{r \in A} \pi(r) \cdot r$$

[1]This reference to a multiset is relevant.

a ***convex combination of the elements of*** A. For instance, given the multiset $A := \{\!\{4, 2, 2\}\!\}$, $\frac{1}{3}4 + \frac{1}{3}2 + \frac{1}{3}2$, so $\frac{8}{3}$, is a convex combination of the elements of A.

To see the use of this notion when discussing mixed strategies note that for every joint mixed strategy m we have

$$p_i(m) = \sum_{s_i \in support(m_i)} m_i(s_i) \cdot p_i(s_i, m_{-i}).$$

That is, $p_i(m)$ is a convex combination of the elements of the multiset

$$\{\!\{p_i(s_i, m_{-i}) \mid s_i \in support(m_i)\}\!\}.$$

We shall employ the following simple observations on convex combinations.

Note 27 (Convex Combination) *Consider a convex combination*

$$cc := \sum_{r \in A} \pi(r) \cdot r$$

of the elements of a finite multiset A of reals. Then

(i) $max\ A \geq cc$,

(ii) $cc \geq max\ A$ *iff*

- $cc = r$ *for all* $r \in A$ *such that* $\pi(r) > 0$,
- $cc \geq r$ *for all* $r \in A$ *such that* $\pi(r) = 0$.

□

Lemma 28 (Characterization) *Consider a finite strategic game*

$$(S_1, \ldots, S_n, p_1, \ldots, p_n).$$

The following statements are equivalent:

(i) m is a Nash equilibrium in mixed strategies, i.e.,

$$p_i(m) \geq p_i(m'_i, m_{-i})$$

for all $i \in \{1, \ldots, n\}$ and all $m'_i \in \Delta S_i$, *(ii)* *for all* $i \in \{1, \ldots, n\}$ *and all* $s_i \in S_i$

$$p_i(m) \geq p_i(s_i, m_{-i}),$$

(iii) *for all* $i \in \{1, \ldots, n\}$ *and all* $s_i \in support(m_i)$

$$p_i(m) = p_i(s_i, m_{-i})$$

and for all $i \in \{1, \ldots, n\}$ *and all* $s_i \notin support(m_i)$

$$p_i(m) \geq p_i(s_i, m_{-i}).$$

Note that the equivalence between (i) and (ii) implies that each Nash equilibrium of the initial game is a pure Nash equilibrium of the mixed extension. In turn, the equivalence between (i) and (iii) provides us with a straightforward way of testing whether a joint mixed strategy is a Nash equilibrium.

Proof.
$(i) \Rightarrow (ii)$ Immediate.
$(ii) \Rightarrow (iii)$ We noticed already that $p_i(m)$ is a convex combination of the elements of the multiset

$$A := \{\!\{p_i(s_i, m_{-i}) \mid s_i \in support(m_i)\}\!\}.$$

So this implication is a consequence of part (ii) of the Convex Combination Note 27.
$(iii) \Rightarrow (i)$ Consider the multiset

$$A := \{\!\{p_i(s_i, m_{-i}) \mid s_i \in S_i\}\!\}.$$

But for all $m'_i \in \Delta S_i$, in particular m_i, the payoff $p_i(m'_i, m_{-i})$ is a convex combination of the elements of the multiset A.

So by the assumptions and part (ii) of the Convex Combination Note 27

$$p_i(m) \geq max\ A,$$

and by part (i) of the above Note

$$max\ A \geq p_i(m'_i, m_{-i}).$$

Hence $p_i(m) \geq p_i(m'_i, m_{-i})$. $\square$

We now illustrate the use of the above theorem by finding in the Battle of the Sexes game a Nash equilibrium in mixed strategies, in addition to the two pure ones exhibited in Chapter 3. Take

$$m_1 := r_1 \cdot F + (1 - r_1) \cdot B,$$
$$m_2 := r_2 \cdot F + (1 - r_2) \cdot B,$$

where $0 < r_1, r_2 < 1$. By definition

$$p_1(m_1, m_2) = 2 \cdot r_1 \cdot r_2 + (1 - r_1) \cdot (1 - r_2),$$
$$p_2(m_1, m_2) = r_1 \cdot r_2 + 2 \cdot (1 - r_1) \cdot (1 - r_2).$$

Suppose now that (m_1, m_2) is a Nash equilibrium in mixed strategies. By the equivalence between (i) and (iii) of the Characterization Lemma 28 $p_1(F, m_2) = p_1(B, m_2)$, i.e., (using $r_1 = 1$ and $r_1 = 0$ in the above formula for $p_1(\cdot)$) $2 \cdot r_2 = 1 - r_2$, and $p_2(m_1, F) = p_2(m_1, B)$, i.e., (using $r_2 = 1$ and $r_2 = 0$ in the above formula for $p_2(\cdot)$) $r_1 = 2 \cdot (1 - r_1)$. So $r_2 = \frac{1}{3}$ and $r_1 = \frac{2}{3}$.

This implies that for these values of r_1 and r_2, (m_1, m_2) is a Nash equilibrium in mixed strategies and we have

$$p_1(m_1, m_2) = p_2(m_1, m_2) = \tfrac{2}{3}.$$

9.3 Nash theorem

We now establish a fundamental result about games that are mixed extensions. In what follows we shall use the following result from the calculus.

Theorem 29 (Extreme Value Theorem) *Suppose that A is a non-empty compact subset of $\mathbb{R}^n$ and*

$$f : A \to \mathbb{R}$$

is a continuous function. Then f attains a minimum and a maximum. $\square$

The example of the Matching Pennies game illustrated that some strategic games do not have a Nash equilibrium. In the case of mixed extensions the situation changes and we have the following fundamental result established by J. Nash in 1950.

Theorem 30 (Nash) *Every mixed extension of a finite strategic game has a Nash equilibrium.*

In other words, every finite strategic game has a Nash equilibrium in mixed strategies. In the case of the Matching Pennies game it is straightforward to check that $(\frac{1}{2} \cdot H + \frac{1}{2} \cdot T, \frac{1}{2} \cdot H + \frac{1}{2} \cdot T)$ is such a Nash equilibrium. In this equilibrium the payoffs to each player are 0.

Nash Theorem follows directly from the following result.[2]

Theorem 31 (Kakutani) *Suppose that A is a non-empty compact and convex subset of $\mathbb{R}^n$ and*

$$\Phi : A \rightarrow \mathcal{P}(A)$$

such that

- *$\Phi(x)$ is non-empty and convex for all $x \in A$,*
- *the **graph** of Φ, so the set $\{(x, y) \mid y \in \Phi(x)\}$, is closed.*

Then $x^* \in A$ exists such that $x^* \in \Phi(x^*)$. □

Proof of Nash Theorem. Fix a finite strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$. Define the function $best_i : \times_{j \neq i} \Delta S_j \rightarrow \mathcal{P}(\Delta S_i)$ by

$$best_i(m_{-i}) := \{m_i \in \Delta S_i \mid m_i \text{ is a best response to } m_{-i}\}.$$

Then define the function $best : \Delta S_1 \times \ldots \times \Delta S_n \rightarrow \mathcal{P}(\Delta S_1 \times \ldots \times \Delta S_n)$ by

$$best(m) := best_1(m_{-1}) \times \ldots \times best_n(m_{-n}).$$

It is now straightforward to check that m is a Nash equilibrium iff $m \in best(m)$. Moreover, one easily can check that the function $best(\cdot)$ satisfies the conditions of Kakutani Theorem. The fact that for every joint mixed strategy m, $best(m)$ is non-empty is a direct consequence of the Extreme Value Theorem 29. □

Ever since Nash established his celebrated Theorem, a search has continued to generalize his result to a larger class of games. A motivation for this endevour has been existence of natural infinite games that are not mixed extensions of finite games. As an example of such an early result let us mention the following theorem stablished independently in 1952 by Debreu, Fan and Glickstein.

[2] Recall that a subset A of $\mathbb{R}^n$ is called ***compact*** if it is closed and bounded.

Theorem 32 *Consider a strategic game such that*

- *each strategy set is a non-empty compact convex subset of* $\mathbb{R}^n$*,*
- *each payoff function* p_i *is continuous and quasi-concave in the ith argument.*[3]

Then a Nash equilibrium exists.

More recent work in this area focused on existence of Nash equilibria in games with non-continuous payoff functions.

9.4 Minimax theorem

Let us return now to strictly competitive games that we studied in Chapter 6. First note the following lemma.

Lemma 33 *Consider a strategic game* $(S_1, \ldots, S_n, p_1, \ldots, p_n)$ *that is a mixed extension. Then*

(i) For all $s_i \in S_i$, $\min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i})$ *exists.*

(ii) $\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i})$ *exists.*

(iii) For all $s_{-i} \in S_{-i}$, $\max_{s_i \in S_i} p_i(s_i, s_{-i})$ *exists.*

(iv) $\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} p_i(s_i, s_{-i})$ *exists.*

Proof. It is a direct consequence of the Extreme Value Theorem 29. □

This lemma implies that we can apply the results of Chapter 6 to each strictly competitive game that is a mixed extension. Indeed, it ensures that the minima and maxima the existence of which we assumed in the proofs given there always exist. However, equipped with the knowledge that each such game has a Nash equilibrium we can now draw additional conclusions.

Theorem 34 *Consider a strictly competitive game that is a mixed extension. For* $i = 1, 2$ *we have* $maxmin_i = minmax_i$.

[3]Recall that the function $p_i : S \to \mathbb{R}$ is ***quasi-concave in the*** i***th argument*** if the set $\{s'_i \in S_i \mid p_i(s'_i, s_{-i}) \geq p_i(s)\}$ is convex for all $s \in S$.

Proof. By the Nash Theorem 30 and the Strictly Competitive Games Theorem 18(ii). $\square$

The formulation 'a strictly competitive game that is a mixed extension' is rather awkward and it is tempting to write instead 'the mixed extension of a strictly competitive game'. However, one can show that the mixed extension of a strictly competitive game does not need to be a strictly competitive game, see Exercise 11.

On the other hand we have the following simple observation.

Note 35 (Mixed Extension) *The mixed extension of a zero-sum game is a zero-sum game.*

Proof. Fix a finite zero-sum game (S_1, S_2, p_1, p_2). For each joint strategy m we have

$$p_1(m)+p_2(m) = \sum_{s\in S} m(s)p_1(s)+\sum_{s\in S} m(s)p_2(s) = \sum_{s\in S} m(s)(p_1(s)+p_2(s)) = 0.$$

$\square$

This means that for finite zero-sum games we have the following result, originally established by von Neumann in 1928.

Theorem 36 (Minimax) *Consider a finite zero-sum game* $G := (S_1, S_2, p_1, p_2)$. *Then for* $i = 1, 2$

$$\max_{m_i\in M_i} \min_{m_{-i}\in M_{-i}} p_i(m_i, m_{-i}) = \min_{m_{-i}\in M_{-i}} \max_{m_i\in M_i} p_i(m_i, m_{-i}).$$

Proof. By the Mixed Extension Note 35 the mixed extension of G is zero-sum, so strictly competitive. It suffices to use Theorem 34 and expand the definitions of $minmax_i$ and $maxmin_i$. $\square$

Finally, note that using the matrix notation we can rewrite the above equalities as follows, where $\mathbf{A}$ is an arbitrary $k \times \ell$ matrix (that is the reward matrix of a zero-sum game):

$$\max_{\mathbf{p}\in\Delta^{k-1}} \min_{\mathbf{q}\in\Delta^{\ell-1}} \mathbf{pAq}^T = \min_{\mathbf{q}\in\Delta^{\ell-1}} \max_{\mathbf{p}\in\Delta^{k-1}} \mathbf{pAq}^T.$$

So the Minimax Theorem can be alternatively viewed as a theorem about matrices and unit simplices. This formulation of the Minimax Theorem has been generalized in many ways to a statement

$$\max_{x \in X} \min_{y \in Y} f(x, y) = \min_{y \in Y} \max_{x \in X} f(x, y),$$

where X and Y are appropriate sets replacing the unit simplices and $f : X \times Y \rightarrow \mathbb{R}$ is an appropriate function replacing the payoff function. Such theorems are called Minimax theorems.

Exercise 11 Find a 2×2 strictly competitive game such that its mixed extension is not a strictly competitive game.

Exercise 12 Prove that the Matching Pennies game has exactly one Nash equilibrium in mixed strategies. $\square$

Chapter 10

Elimination by Mixed Strategies

The notions of dominance apply in particular to mixed extensions of finite strategic games. But we can also consider dominance of a *pure* strategy by a *mixed* strategy. Given a finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$, we say that a (pure) strategy s_i of player i is ***strictly dominated by*** a mixed strategy m_i if

$$\forall s_{-i} \in S_{-i} \; p_i(m_i, s_{-i}) > p_i(s_i, s_{-i}),$$

and that s_i is ***weakly dominated by*** a mixed strategy m_i if

$$\forall s_{-i} \in S_{-i} \; p_i(m_i, s_{-i}) \geq p_i(s_i, s_{-i}) \text{ and } \exists s_{-i} \in S_{-i} \; p_i(m_i, s_{-i}) > p_i(s_i, s_{-i}).$$

In what follows we discuss for these two forms of dominance the counterparts of the results presented in Chapters 3 and 4.

10.1 Elimination of strictly dominated strategies

Strict dominance by a mixed strategy leads to a stronger form of strategy elimination. For example, in the game

	L	R
T	$2, 1$	$0, 1$
M	$0, 1$	$2, 1$
B	$0, 1$	$0, 1$

the strategy B is strictly dominated neither by T nor M but is strictly dominated by $\frac{1}{2} \cdot T + \frac{1}{2} \cdot M$.

We now focus on iterated elimination of pure strategies that are strictly dominated by a mixed strategy. As in Chapter 3 we would like to clarify whether it affects the Nash equilibria, in this case equilibria in mixed strategies. We denote the corresponding reduction relation between restrictions of a finite strategic game by $\rightarrow_{SM}$.

First, we introduce the following notation. Given two mixed strategies m_i, m'_i and a strategy s_i we denote by $m_i[s_i/m'_i]$ the mixed strategy obtained from m_i by substituting the strategy s_i by m'_i and by 'normalizing' the resulting sum. For example, given $m_i = \frac{1}{3}H + \frac{2}{3}T$ and $m'_i = \frac{1}{2}H + \frac{1}{2}T$ we have $m_i[H/m'_i] = \frac{1}{3}(\frac{1}{2}H + \frac{1}{2}T) + \frac{2}{3}T = \frac{1}{6}H + \frac{5}{6}T$.

We also use the following identification of mixed strategies over two sets of strategies S'_i and S_i such that $S'_i \subseteq S_i$. We view a mixed strategy $m_i \in \Delta S_i$ such that $support(m_i) \subseteq S'_i$ as a mixed strategy 'over' the set S'_i, i.e., as an element of $\Delta S'_i$, by limiting the domain of m_i to S'_i. Further, we view each mixed strategy $m_i \in \Delta S'_i$ as a mixed strategy 'over' the set S_i, i.e., as an element of ΔS_i, by assigning the probability 0 to the elements in $S_i \setminus S'_i$.

Next, we establish the following auxiliary lemma.

Lemma 37 (Persistence) *Given a finite strategic game G consider two restrictions R and R' of G such that $R \rightarrow_{SM} R'$.*

Suppose that a strategy $s_i \in R_i$ is strictly dominated in R by a mixed strategy from R. Then s_i is strictly dominated in R by a mixed strategy from R'.

Proof. We shall use the following, easy to establish, two properties of strict dominance by a mixed strategy in a given restriction:

(a) for all $\alpha \in (0, 1]$, if s_i is strictly dominated by $(1 - \alpha)s_i + \alpha\, m_i$, then s_i is strictly dominated by m_i,

(b) if s_i is strictly dominated by m_i and s'_i is strictly dominated by m'_i, then s_i is strictly dominated by $m_i[s'_i/m'_i]$.

Suppose that $R_i \setminus R'_i = \{t^1_i, \ldots, t^k_i\}$. By definition for all $j \in \{1, \ldots, k\}$ there exists in R a mixed strategy m^j_i such that t^j_i is strictly dominated in R

by m_i^j. We first prove by complete induction that for all $j \in \{1, \ldots, k\}$ there exists in R a mixed strategy n_i^j such that

$$t_i^j \text{ is strictly dominated in } R \text{ by } n_i^j \text{ and } \mathit{support}(n_i^j) \cap \{t_i^1, \ldots, t_i^j\} = \emptyset. \tag{10.1}$$

For some $\alpha \in (0, 1]$ and a mixed strategy n_i^1 with $t_i^1 \notin \mathit{support}(n_i^1)$ we have

$$m_i^1 = (1 - \alpha)t_i^1 + \alpha\, n_i^1.$$

By assumption t_i^1 is strictly dominated in R by m_i^1, so by (a) t_i^1 is strictly dominated in R by n_i^1, which proves (10.1) for $j = 1$.

Assume now that $\ell < k$ and that (10.1) holds for all $j \in \{1, \ldots, \ell\}$. By assumption $t_i^{\ell+1}$ is strictly dominated in R by $m_i^{\ell+1}$.

Let

$$m_i'' := m_i^{\ell+1}[t_i^1/n_i^1]\ldots[t_i^\ell/n_i^\ell].$$

By the induction hypothesis and (b) $t_i^{\ell+1}$ is strictly dominated in R by m_i'' and $\mathit{support}(m_i'') \cap \{t_i^1, \ldots, t_i^\ell\} = \emptyset$.

For some $\alpha \in (0, 1]$ and a mixed strategy $n_i^{\ell+1}$ with $t_i^{\ell+1} \notin \mathit{support}(n_i^{\ell+1})$ we have

$$m_i'' = (1 - \alpha)t_i^{\ell+1} + \alpha\, n_i^{\ell+1}.$$

By (a) $t_i^{\ell+1}$ is strictly dominated in R by $n_i^{\ell+1}$. Also $\mathit{support}(n_i^{\ell+1}) \cap \{t_i^1, \ldots, t_i^{\ell+1}\} = \emptyset$, which proves (10.1) for $j = \ell + 1$.

Suppose now that the strategy s_i is strictly dominated in R by a mixed strategy m_i from R. Define

$$m_i' := m_i[t_i^1/n_i^1]\ldots[t_i^k/n_i^k].$$

Then by (b) and (10.1) s_i is strictly dominated in R by m_i' and $\mathit{support}(m_i') \subseteq R_i'$, i.e., m_i' is a mixed strategy in R'. □

The following is a counterpart of the Strict Elimination Lemma 1 and will be used in a moment.

Lemma 38 (Strict Mixed Elimination) *Given a finite strategic game G consider two restrictions R and R' of G such that $R \rightarrow_{SM} R'$.*

Then m is a Nash equilibrium of R iff it is a Nash equilibrium of R'.

Proof. Let

$$R := (R_1, \ldots, R_n, p_1, \ldots, p_n),$$

and

$$R' := (R'_1, \ldots, R'_n, p_1, \ldots, p_n).$$

($\Rightarrow$) It suffices to show that m is also a joint mixed strategy in R', i.e., that for all $i \in \{1, \ldots, n\}$ we have $support(m_i) \subseteq R'_i$.

Suppose otherwise. Then for some $i \in \{1, \ldots, n\}$ a strategy $s_i \in support(m_i)$ is strictly dominated by a mixed strategy $m'_i \in \Delta R_i$. So

$$p_i(m'_i, m''_{-i}) > p_i(s_i, m''_{-i}) \text{ for all } m''_{-i} \in \times_{j \neq i} \Delta R_j.$$

In particular

$$p_i(m'_i, m_{-i}) > p_i(s_i, m_{-i}).$$

But m is a Nash equilibrium of R and $s_i \in support(m_i)$ so by the Characterization Lemma 28

$$p_i(m) = p_i(s_i, m_{-i}).$$

Hence

$$p_i(m'_i, m_{-i}) > p_i(m),$$

which contradicts the choice of m.

($\Leftarrow$) Suppose m is not a Nash equilibrium of R. Then by the Characterization Lemma 28 for some $i \in \{1, \ldots, n\}$ and $s'_i \in R_i$

$$p_i(s'_i, m_{-i}) > p_i(m).$$

The strategy s'_i is eliminated since m is a Nash equilibrium of R'. So s'_i is strictly dominated in R by some mixed strategy in R. By the Persistence Lemma 37 s'_i is strictly dominated in R by some mixed strategy m'_i in R'. So

$$p_i(m'_i, m''_{-i}) \geq p_i(s'_i, m''_{-i}) \text{ for all } m''_{-i} \in \times_{j \neq i} \Delta R_j.$$

In particular

$$p_i(m'_i, m_{-i}) \geq p_i(s'_i, m_{-i})$$

and hence by the choice of s'_i

$$p_i(m'_i, m_{-i}) > p_i(m).$$

Since $m'_i \in \Delta R'_i$ this contradicts the assumption that m is a Nash equilibrium of R'. $\Box$

Instead of the lengthy wording 'the iterated elimination of strategies strictly dominated by a mixed strategy' we write ***IESDMS***. We have then the following counterpart of the IESDS Theorem 2, where we refer to Nash equilibria in mixed strategies. Given a restriction G' of G and a joint mixed strategy m of G, when we say that m is a Nash equilibrium of G' we implicitly stipulate that all supports of all m_is consist of strategies from G'.

We then have the following counterpart of the IESDS Theorem 2.

Theorem 39 (IESDMS) *Suppose that G is a finite strategic game.*

(i) If G' is an outcome of IESDMS from G, then m is a Nash equilibrium of G iff it is a Nash equilibrium of G'.

(ii) If G is solved by IESDMS, then the resulting joint strategy is a unique Nash equilibrium of G (in, possibly, mixed strategies).

Proof. By the Strict Mixed Elimination Lemma 38. $\Box$

To illustrate the use of this result let us return to the beauty contest game discussed in Examples 2 of Chapter 1 and 10 in Chapter 4. We explained there that $(1, \ldots, 1)$ is *a* Nash equilibrium. Now we can draw a stronger conclusion.

Example 19 One can show that the beauty contest game is solved by IESDMS in 99 rounds. In each round the highest strategy of each player is removed and eventually each player is left with the strategy 1. On the account of the above theorem we now conclude that $(1, \ldots, 1)$ is a *unique* Nash equilibrium. $\Box$

As in the case of strict dominance by a pure strategy we now address the question whether the outcome of IESDMS is unique. The answer is positive. To establish this result we proceed as before and establish the following lemma first. Recall that the notion of hereditarity was defined in the Appendix of Chapter 3.

Lemma 40 (Hereditarity III) *The relation of being strictly dominated by a mixed strategy is hereditary on the set of restrictions of a given finite game.*

Proof. This is an immediate consequence of the Persistence Lemma 37. Indeed, consider a finite strategic game G and two restrictions R and R' of G such that $R \rightarrow_{SM} R'$.

Suppose that a strategy $s_i \in R'_i$ is strictly dominated in R by a mixed strategy in R. By the Persistence Lemma 37 s_i is strictly dominated in R by a mixed strategy in R'. So s_i is also strictly dominated in R' by a mixed strategy in R'. $\Box$

This brings us to the following conclusion.

Theorem 41 (Order independence III) *All iterated eliminations of strategies strictly dominated by a mixed strategy yield the same outcome.*

Proof. By Theorem 5 and the Hereditarity III Lemma 40. $\Box$

10.2 Elimination of weakly dominated strategies

Next, we consider iterated elimination of pure strategies that are weakly dominated by a mixed strategy.

As already noticed in Chapter 4 an elimination by means of weakly dominated strategies can result in a loss of Nash equilibria. Clearly, the same observation applies here. On the other hand, as in the case of pure strategies, we can establish a partial result, where we refer to the reduction relation $\rightarrow_{WM}$ with the expected meaning.

Lemma 42 (Mixed Weak Elimination) *Given a finite strategic game G consider two restrictions R and R' of G such that $R \rightarrow_{WM} R'$.*

If m is a Nash equilibrium of R', then it is a Nash equilibrium of R.

Proof. It suffices to note that both the proofs of the Persistence Lemma 37 and of the ($\Leftarrow$) implication of the Strict Mixed Elimination Lemma 38 apply without any changes to weak dominance, as well. $\Box$

This brings us to the following counterpart of the IEWDS Theorem 9, where we refer to Nash equilibria in mixed strategies. Instead of 'the iterated elimination of strategies weakly dominated by a mixed strategy' we write ***IEWDMS***.

Theorem 43 (IEWDMS) *Suppose that G is a finite strategic game.*

(i) If G' is an outcome of IEWDMS from G and m is a Nash equilibrium of G', then m is a Nash equilibrium of G.

(ii) If G is solved by IEWDMS, then the resulting joint strategy is a Nash equilibrium of G.

Proof. By the Mixed Weak Elimination Lemma 42. □

Here is a simple application of this theorem.

Corollary 44 *Every mixed extension of a finite strategic game has a Nash equilibrium such that no strategy used in it is weakly dominated by a mixed strategy.*

Proof. It suffices to apply Nash Theorem 30 to an outcome of IEWDMS and use item (i) of the above theorem. □

Finally, observe that the outcome of IEWMDS does not need to be unique. In fact, Example 11 applies here, as well. It is instructive to note where the proof of the Order independence III Theorem 41 breaks down. It happens in the very last step of the proof of the Hereditarity III Lemma 40. Namely, if $R \rightarrow_{WM} R'$ and a strategy $s_i \in R'_i$ is weakly dominated in R by a mixed strategy in R', then we cannot conclude that s_i is weakly dominated in R' by a mixed strategy in R'.

10.3 Rationalizability

Finally, we consider iterated elimination of strategies that are never best responses to a joint mixed strategy of the opponents. Strategies that survive such an elimination process are called rationalizable strategies.

Formally, we define rationalizable strategies as follows. Consider a restriction R of a finite strategic game G. Let

$$\mathcal{RAT}(R) := (S'_1, \ldots, S'_n),$$

where for all $i \in \{1, \ldots, n\}$

$$S'_i := \{s_i \in R_i \mid \exists m_{-i} \in \times_{j \neq i} \Delta R_j \ s_i \text{ is a best response to } m_{-i} \text{ in } G\}.$$

Note the use of G instead of R in the definition of S'_i. We shall comment on it in below.

Consider now the outcome $G_{\mathcal{RAT}}$ of iterating $\mathcal{RAT}$ starting with G. We call then the strategies present in the restriction $G_{\mathcal{RAT}}$ ***rationalizable***.

We have the following counterpart of the IESDMS Theorem 39.

Theorem 45 *Assume a finite strategic game G.*

(i) Then m is a Nash equilibrium of G iff it is a Nash equilibrium of $G_{\mathcal{RAT}}$.

(ii) If each player has in $G_{\mathcal{RAT}}$ exactly one strategy, then the resulting joint strategy is a unique Nash equilibrium of G.

In the context of rationalizability a joint mixed strategy of the opponents is referred to as a ***belief***. The definition of rationalizability is generic in the class of beliefs w.r.t. which best responses are collected. For example, we could use here joint pure strategies of the opponents, or probability distributions over the Cartesian product of the opponents' strategy sets, so the elements of the set ΔS_{-i} (extending in an expected way the payoff functions). In the first case we talk about ***point beliefs*** and in the second case about ***correlated beliefs***.

In the case of point beliefs we can apply the elimination procedure entailed by $\mathcal{RAT}$ to arbitrary games. To avoid discussion of the outcomes reached in the case of infinite iterations we focus on a result for a limited case. We refer here to Nash equilibria in pure strategies.

Theorem 46 *Assume a strategic game G. Consider the definition of the $\mathcal{RAT}$ operator for the case of point beliefs and suppose that the outcome $G_{\mathcal{RAT}}$ is reached in finitely many steps.*

(i) Then s is a Nash equilibrium of G iff it is a Nash equilibrium of $G_{\mathcal{RAT}}$.

(ii) If each player is left in $G_{\mathcal{RAT}}$ with exactly one strategy, then the resulting joint strategy is a unique Nash equilibrium of G.

A subtle point is that when G is infinite, the restriction $G_{\mathcal{RAT}}$ may have empty strategy sets (and hence no joint strategy).

Example 20 ***Bertrand competition*** is a game concerned with a simultaneous selection of prices for the same product by two firms. The product is then sold by the firm that chose a lower price. In the case of a tie the product is sold by both firms and the profits are split.

Consider a version in which the range of possible prices is the left-open real interval $(0, 100]$ and the demand equals $100 - p$, where p is the lower price. So in this game G there are two players, each with the set $(0, 100]$ of strategies and the payoff functions are defined by:

$$p_1(s_1, s_2) := \begin{cases} s_1(100 - s_1) & \text{if } s_1 < s_2 \\ \dfrac{s_1(100 - s_1)}{2} & \text{if } s_1 = s_2 \\ 0 & \text{if } s_1 > s_2 \end{cases}$$

$$p_2(s_1, s_2) := \begin{cases} s_2(100 - s_2) & \text{if } s_2 < s_1 \\ \dfrac{s_2(100 - s_2)}{2} & \text{if } s_1 = s_2 \\ 0 & \text{if } s_2 > s_1 \end{cases}$$

Consider now each player's best responses to the strategies of the opponent. Since $s_1 = 50$ maximizes the value of $s_1(100 - s_1)$ in the interval $(0, 100]$, the strategy 50 is the unique best response of the first player to any strategy $s_2 > 50$ of the second player. Further, no strategy is a best response to a strategy $s_2 \leq 50$. By symmetry the same holds for the strategies of the second player.

So the elimination of never best responses leaves each player with a single strategy, 50. In the second round we need to consider the best responses to these two strategies in the *original* game G. In G the strategy $s_1 = 49$ is a better response to $s_2 = 50$ than $s_1 = 50$ and symmetrically for the second player. So in the second round of elimination both strategies 50 are eliminated and we reach the restriction with the empty strategy sets. By Theorem 46 we conclude that the original game G has no Nash equilibrium. □

Note that if we defined S'_i in the definition of the operator $\mathcal{RAT}$ using the restriction R instead of the original game G, the iteration would stop in the above example after the first round. Such a modified definition of the $\mathcal{RAT}$ operator is actually an instance of the IENBR (iterated elimination of never best responses) in which at each stage all never best responses are eliminated. So for the above game G we can then conclude by the IENBR

Theorem 11(i) that it has at most one equilibrium, namely $(50, 50)$, and then check separately that in fact it is not a Nash equilibrium.

Exercise 13 Show that the beauty contest game is indeed solved by IESDMS in 99 rounds. □

10.4 A comparison between the introduced notions

We introduced so far the notions of strict dominance, weak dominance, and a best response, and related them to the notion of a Nash equilibrium. To conclude this section we clarify the connections between the notions of dominance and of best response.

Clearly, if a strategy is strictly dominated, then it is a never best response. However, the converse fails. Further, there is no relation between the notions of weak dominance and never best response. Indeed, in the game considered in Section 4.2 strategy C is a never best response, yet it is neither strictly nor weakly dominated. Further, in the game given in Example 11 strategy M is weakly dominated and is also a best response to B.

The situation changes in the case of mixed extensions of two-player finite games. Below by a ***totally mixed strategy*** we mean a mixed strategy with full support, i.e., one in which each strategy is used with a strictly positive probability. We have the following results.

Theorem 47 *Consider a finite two-player strategic game.*

(i) A pure strategy is strictly dominated by a mixed strategy iff it is not a best response to a mixed strategy.

(ii) A pure strategy is weakly dominated by a mixed strategy iff it is not a best response to a totally mixed strategy.

We only prove here part (i).

We shall use the following result.

Theorem 48 (Separating Hyperplane) *Let A and B be disjoint convex subsets of $\mathbb{R}^k$. Then there exists a nonzero $c \in \mathbb{R}^k$ and $d \in \mathbb{R}$ such that*

$$c \cdot x \geq d \text{ for all } x \in A,$$

$$c \cdot y \leq d \text{ for all } y \in B.$$

Proof of Theorem 47(i).

Clearly, if a pure strategy is strictly dominated by a mixed strategy, then it is not a best response to a mixed strategy. To prove the converse fix a two-player strategic game (S_1, S_2, p_1, p_2). Also fix $i \in \{1, 2\}$.

Suppose that a strategy $s_i \in S_i$ is not strictly dominated by a mixed strategy. Let

$$A := \{x \in \mathbb{R}^{|S_{-i}|} \mid \forall s_{-i} \in S_{-i} \ x_{s_{-i}} > 0\}$$

and

$$B := \{(p_i(m_i, s_{-i}) - p_i(s_i, s_{-i}))_{s_{-i} \in S_{-i}} \mid m_i \in \Delta S_i\}.$$

By the choice of s_i the sets A and B are disjoint. Moreover, both sets are convex subsets of $\mathbb{R}^{|S_{-i}|}$.

By the Separating Hyperplane Theorem 48 for some nonzero $c \in \mathbb{R}^{|S_{-i}|}$ and $d \in \mathbb{R}$

$$c \cdot x \geq d \text{ for all } x \in A, \tag{10.2}$$

$$c \cdot y \leq d \text{ for all } y \in B. \tag{10.3}$$

But $\mathbf{0} \in B$, so by (10.3) $d \geq 0$. Hence by (10.2) and the definition of A for all $s_{-i} \in S_{-i}$ we have $c_{s_{-i}} \geq 0$. Again by (10.2) and the definition of A this excludes the contingency that $d > 0$, i.e., $d = 0$. Hence by (10.3)

$$\textstyle\sum_{s_{-i} \in S_{-i}} c_{s_{-i}} p_i(m_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} c_{s_{-i}} p_i(s_i, s_{-i}) \text{ for all } m_i \in \Delta S_i. \tag{10.4}$$

Let $\bar{c} := \sum_{s_{-i} \in S_{-i}} c_{s_{-i}}$. By the assumption $\bar{c} \neq 0$. Take

$$m_{-i} := \sum_{s_{-i} \in S_{-i}} \frac{c_{s_{-i}}}{\bar{c}} s_{-i}.$$

Then (10.4) can be rewritten as

$$p_i(m_i, m_{-i}) \leq p_i(s_i, m_{-i}) \text{ for all } m_i \in \Delta S_i,$$

i.e., s_i is a best response to m_{-i}. □

Chapter 11

Mechanism Design

Mechanism design is one of the important areas of economics. The 2007 Nobel prize in Economics went to three economists who laid its foundations. To quote from the article *Intelligent design*, published in *The Economist*, October 18th, 2007, mechanism design deals with the problem of 'how to arrange our economic interactions so that, when everyone behaves in a self-interested manner, the result is something we all like.' So these interactions are supposed to yield desired social decisions when each agent is interested in maximizing only his own utility.

In mechanism design one is interested in the ways of inducing the players to submit true information. To discuss it in more detail we need to introduce some basic concepts.

11.1 Decision problems

Assume a set of ***decisions*** D, a set $\{1, \ldots, n\}$ of players, and for each player

- a set of ***types*** Θ_i, and
- an ***initial utility function*** $v_i : D \times \Theta_i \to \mathbb{R}$.

In this context a type is some private information known only to the player, for example, in the case of an auction, player's valuation of the items for sale.

As in the case of strategy sets we use the following abbreviations:

- $\Theta := \Theta_1 \times \ldots \times \Theta_n$,

- $\Theta_{-i} := \Theta_1 \times \ldots \times \Theta_{i-1} \times \Theta_{i+1} \times \ldots \times \Theta_n$, and similarly with θ_{-i} where $\theta \in \Theta$,
- $(\theta'_i, \theta_{-i}) := \theta_1 \times \ldots \times \theta_{i-1} \times \theta'_i \times \theta_{i+1} \times \ldots \times \theta_n$.

In particular $(\theta_i, \theta_{-i}) = \theta$.

A ***decision rule*** is a function $f : \Theta \to D$. We call the tuple

$$(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$$

a ***decision problem***.

Decision problems are considered in the presence of a ***central authority*** who takes decisions on the basis of the information provided by the players. Given a decision problem the desired decision is obtained through the following sequence of events, where f is a given, publicly known, decision rule:

- each player i receives (becomes aware of) his type $\theta_i \in \Theta_i$,
- each player i announces to the central authority a type $\theta'_i \in \Theta_i$; this yields a type vector $\theta' := (\theta'_1, \ldots, \theta'_n)$,
- the central authority then takes the decision $d := f(\theta')$ and communicates it to each player,
- the resulting initial utility for player i is then $v_i(d, \theta_i)$.

The difficulty in taking decisions through the above described sequence of events is that players are assumed to be ***rational***, that is they want to maximize their utility. As a result they may submit false information to manipulate the outcome (decision). We shall return to this problem in the next section. But first, to better understand the above notion let us consider some natural examples.

For a sequence θ of reals we shall here use the notation $\text{argsmax}\theta$ originally introduced in Chapter 7. Additionally, for a function $g : A \to \mathbb{R}$ we define

$$\text{argmax}_{x \in A} g(x) := \{y \in A \mid g(y) = \max_{x \in A} g(x)\}.$$

So $a \in \text{argmax}_{x \in A} g(x)$ means that a is a maximum of the function g on the set A.

Example 21 [**Sealed-bid auction**]

We consider the sealed-bid auction, originally considered in Chapter 7. The important difference is that now we can dispense with the assumption that players' valuations are publicly known by viewing each player's valuation as his type. More precisely, we model this type of auction as the following decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$:

- $D = \{1, \ldots, n\}$,
- for all $i \in \{1, \ldots, n\}$, $\Theta_i = \mathbb{R}_+$;

  $\theta_i \in \Theta_i$ is player's i valuation of the object,
- $v_i(d, \theta_i) := \begin{cases} \theta_i & \text{if } d = i \\ 0 & \text{otherwise} \end{cases}$
- $f(\theta) := \text{argsmax}\, \theta$.

  Here decision $d \in D$ indicates to which player the object is sold. Note that at this stage we only modeled the fact that the object is sold to the highest bidder (with the ties resolved in the favour of a bidder with the lowest index). We shall return to the problem of payments in the next section. □

Example 22 [**Public project problem**]

This problem deals with the task of taking a joint decision concerning construction of a ***public good***[1], for example a bridge.

It is explained as follows in the *Scientific Background* of the Royal Swedish Academy of Sciences Press Release that accompanied the Nobel prize in Economics in 2007:

> Each person is asked to report his or her willingness to pay for the project, and the project is undertaken if and only if the aggregate reported willingness to pay exceeds the cost of the project.

[1] In Economics public goods are so-called not excludable and nonrival goods. To quote from the book *N.G. Mankiw, Principles of Economics*, 2nd Editiona, Harcourt, 2001: "People cannot be prevented from using a public good, and one person's enjoyment of a public good does not reduce another person's enjoyment of it."

So there are two decisions: to carry out the project or not. In the terminology of the decision problems each player reports to the central authority his appreciation of the gain from the project when it takes place. If the sum of the appreciations exceeds the cost of the project, the project takes place. We assume that each player has to pay then the same fraction of the cost. Otherwise the project is cancelled.

This leads to the following decision problem:

- $D = \{0, 1\}$,
- each Θ_i is $\mathbb{R}_+$,
- $v_i(d, \theta_i) := d(\theta_i - \frac{c}{n})$,
- $f(\theta) := \begin{cases} 1 & \text{if } \sum_{i=1}^{n} \theta_i \geq c \\ 0 & \text{otherwise} \end{cases}$

Here c is the cost of the project. If the project takes place ($d = 1$), $\frac{c}{n}$ is the cost share of the project for each player. □

Example 23 [Taking an efficient decision] We assume a finite set of decisions. Each player submits to the central authority a function that describes his satisfaction level from each decision if it is taken. The central authority then chooses a decision that yields the maximal overall satisfaction.

This problem corresponds to the following decision problem:

- D is the given finite set of decisions,
- each Θ_i is $\{f \mid f : D \to \mathbb{R}\}$,
- $v_i(d, \theta_i) := \theta_i(d)$,
- the decision rule f is a function such that for all θ, $f(\theta) \in \text{argmax}_{d \in D} \sum_{i=1}^{n} \theta_i(d)$.

□

Example 24 [Reversed sealed-bid auction]

In the ***reversed sealed-bid auction*** each player offers the same service, for example to construct a bridge. The decision is taken by means of a sealed-bid auction. Each player simultaneously submits to the central authority his type (bid) in a sealed envelope and the service is purchased from the lowest bidder.

We model it in exactly the same way as the sealed-bid auction, with the only exception that for each player the types are now non-positive reals. So we consider the following decision problem:

- $D = \{1, \ldots, n\}$,
- for all $i \in \{1, \ldots, n\}$, $\Theta_i = \mathbb{R}_-$ (the set of non-positive reals);

  $-\theta_i$, where $\theta_i \in \Theta_i$, is player's i offer for the service,
- $v_i(d, \theta_i) := \begin{cases} \theta_i & \text{if } d = i \\ 0 & \text{otherwise} \end{cases}$
- $f(\theta) := \text{argsmax}\,\theta$.

Here decision $d \in D$ indicates from which player the service is bought. So for example $f(-8, -5, -4, -6) = 3$, that is, given the offers $8, 5, 4, 6$ (in that order), the service is bought from player 3, since he submitted the lowest bid, namely 4. As in the case of the sealed-bid auction, we shall return to the problem of payments in the next section. □

Example 25 [Buying a path in a network]

We consider a communication network, modelled as a directed graph $G := (V, E)$ (with no self-cycles or parallel edges). We assume that each edge $e \in E$ is owned by a player, also denoted by e. So different edges are owned by different players. We fix two distinguished vertices $s, t \in V$. Each player e submits the cost θ_e of using the edge e. The central authority selects on the basis of players' submissions the shortest $s - t$ path in G.

Below we denote by $G(\theta)$ the graph G augmented with the costs of edges as specified by θ. That is, the cost of each edge i in $G(\theta)$ is θ_i.

This problem can be modelled as the following decision problem:

- $D = \{p \mid p \text{ is a } s - t \text{ path in } G\}$,
- each Θ_i is $\mathbb{R}_+$;

  θ_i is the cost incurred by player i if the edge i is used in the selected path,

- $v_i(p, \theta_i) := \begin{cases} -\theta_i & \text{if } i \in p \\ 0 & \text{otherwise} \end{cases}$

- $f(\theta) := p$, where p is the shortest $s - t$ path in $G(\theta)$.

  In the case of multiple shortest paths we select, say, the one that is alphabetically first.

Note that in the case an edge is selected, the utility of its owner becomes negative. This reflects the fact we focus on incurring costs and not on benefits. In the next section we shall introduce taxes and discuss a scheme according to which each owner of a selected path is *paid* by the central authority an amount exceeding the incurred costs. □

Let us return now to the decision rules. We call a decision rule f ***efficient*** if for all $\theta \in \Theta$ and $d' \in D$

$$\sum_{i=1}^{n} v_i(f(\theta), \theta_i) \geq \sum_{i=1}^{n} v_i(d', \theta_i),$$

or alternatively

$$f(\theta) \in \text{argmax}_{d \in D} \sum_{i=1}^{n} v_i(d, \theta_i).$$

This means that for all $\theta \in \Theta$, $f(\theta)$ is a decision that maximizes the ***initial social welfare***, defined by $\sum_{i=1}^{n} v_i(d, \theta_i)$.

It is easy to check that the decision rules used in Examples 21–25 are efficient. Take for instance Example 25. For each $s - t$ path p we have $\sum_{i=1}^{n} v_i(p, \theta_i) = -\sum_{j \in p} \theta_j$, so $\sum_{i=1}^{n} v_i(p, \theta_i)$ reaches maximum when p is a shortest $s - t$ path in $G(\theta)$, which is the choice made by the decision rule f used there.

11.2 Direct mechanisms

Let us return now to the subject of manipulations. A problem with our description of the sealed-bid auction is that we intentionally neglected the fact that the winner should pay for the object for sale. Still, we can imagine in this limited setting that player i with a strictly positive valuation of the object somehow became aware of the types (that is, bids) of the other players.

Then he should just submit a type strictly larger than the other types. This way the object will be allocated to him and his utility will increase from 0 to θ_i.

The manipulations are more natural to envisage in the case of the public project problem. A player whose type (that is, appreciation of the gain from the project) exceeds $\frac{c}{n}$, the cost share he is to pay, should manipulate the outcome and announce the type c. This will guarantee that the project will take place, irrespectively of the types announced by the other players. Analogously, player whose type is lower than $\frac{c}{n}$ should submit the type 0 to minimize the chance that the project will take place.

To prevent such manipulations we use taxes. This leads to mechanisms that are constructed by combining decision rules with taxes (transfer payments). Each such mechanism is obtained by modifying the initial decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ to the following one:

- the set of decisions is $D \times \mathbb{R}^n$,
- the decision rule is a function $(f, t) : \Theta \to D \times \mathbb{R}^n$, where $t : \Theta \to \mathbb{R}^n$ and $(f, t)(\theta) := (f(\theta), t(\theta))$,
- the ***final utility function*** for player i is a function $u_i : D \times \mathbb{R}^n \times \Theta_i \to \mathbb{R}$ defined by

  $$u_i(d, t_1, \ldots, t_n, \theta_i) := v_i(d, \theta_i) + t_i.$$

  (So defined utilities are called ***quasilinear***.)

We call $(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f, t))$ a ***direct mechanism*** and refer to t as the ***tax function***.

So when the received (true) type of player i is θ_i and his announced type is θ'_i, his final utility is

$$u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i) = v_i(f(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i}),$$

where θ_{-i} are the types announced by the other players.

In each direct mechanism, given the vector θ of announced types, $t(\theta) := (t_1(\theta), \ldots, t_n(\theta))$ is the vector of the resulting payments that the players have to make. If $t_i(\theta) \geq 0$, player i ***receives*** from the central authority $t_i(\theta)$, and if $t_i(\theta) < 0$, he ***pays*** to the central authority $|t_i(\theta)|$.

The following definition then captures the idea that taxes prevent manipulations. We say that a direct mechanism with tax function t is ***incentive compatible*** if for all $\theta \in \Theta$, $i \in \{1, \ldots, n\}$ and $\theta_i' \in \Theta_i$

$$u_i((f,t)(\theta_i, \theta_{-i}), \theta_i) \geq u_i((f,t)(\theta_i', \theta_{-i}), \theta_i).$$

Intuitively, this means that announcing one's true type (θ_i) is better than announcing another type (θ_i'). That is, false announcements, i.e., manipulations do not pay off.

From now on we focus on specific incentive compatible direct mechanisms. Each ***Groves mechanism*** is a direct mechanism obtained by using a tax function $t := (t_1, \ldots, t_n)$, where for all $i \in \{1, \ldots, n\}$

- $t_i : \Theta \rightarrow \mathbb{R}$ is defined by $t_i(\theta) := g_i(\theta) + h_i(\theta_{-i})$, where
- $g_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j)$,
- $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ is an arbitrary function.

Note that $v_i(f(\theta), \theta_i) + g_i(\theta) = \sum_{j=1}^{n} v_j(f(\theta), \theta_j)$ is simply the initial social welfare from the decision $f(\theta)$. In this context the ***final social welfare*** is defined as $\sum_{i=1}^{n} u_i((f,t)(\theta), \theta_i)$, so it equals the sum of the initial social welfare and all the taxes.

The importance of Groves mechanisms is then revealed by the following crucial result due to T. Groves.

Theorem 49 *Consider a decision problem* $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ *with an efficient decision rule* f. *Then each Groves mechanism is incentive compatible.*

Proof. The proof is remarkably straightforward. Since f is efficient, for all $\theta \in \Theta$, $i \in \{1, \ldots, n\}$ and $\theta_i' \in \Theta_i$ we have

$$\begin{aligned} u_i((f,t)(\theta_i, \theta_{-i}), \theta_i) &= \sum_{j=1}^{n} v_j(f(\theta_i, \theta_{-i}), \theta_j) + h_i(\theta_{-i}) \\ &\geq \sum_{j=1}^{n} v_j(f(\theta_i', \theta_{-i}), \theta_j) + h_i(\theta_{-i}) \\ &= u_i((f,t)(\theta_i', \theta_{-i}), \theta_i). \end{aligned}$$

$\square$

When for a given direct mechanism for all θ' we have $\sum_{i=1}^{n} t_i(\theta') \leq 0$, the mechanism is called ***feasible*** (which means that it can be realized without external financing) and when for all θ' we have $\sum_{i=1}^{n} t_i(\theta') = 0$, the mechanism is called ***budget balanced*** (which means that it can be realized without a deficit).

Each Groves mechanism is uniquely determined by the functions $h_1, \ldots,$ h_n. A special case, called ***pivotal mechanism*** is obtained by using

$$h_i(\theta_{-i}) := -\max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

So then

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

Hence for all θ and $i \in \{1, \ldots, n\}$ we have $t_i(\theta) \leq 0$, which means that each player needs to make the payment $|t_i(\theta)|$ to the central authority. In particular, the pivotal mechanism is feasible.

11.3 Back to our examples

When applying Theorem 49 to a specific decision problem we need first to check that the used decision rule is efficient. We noted already that this is the case in Examples 21–25. So in each example Theorem 49 applies and in particular the pivotal mechanism can be used. Let us see now the details of this and other Groves mechanisms for these examples.

Sealed-bid auction

To compute the taxes we use the following observation.

Note 50 *In the sealed-bid auction we have for the pivotal mechanism*

$$t_i(\theta) = \begin{cases} -\max_{j \neq i} \theta_j & \text{if } i = argsmax\, \theta. \\ 0 & \text{otherwise} \end{cases}$$

$\square$

So the highest bidder wins the object and pays for it the amount $\max_{j \neq i} \theta_j$, i.e., the second highest bid. This shows that the pivotal mechanism for the sealed-bid auction is simply the second-price auction proposed by W. Vickrey. By the above considerations this auction is incentive compatible.

In contrast, the first-price sealed-bid auction, in which the winner pays the price he offered, is not incentive compatible. Indeed, suppose that the true types are (4,5,7) and that players 1 and 2 bid truthfully. If player 3 bids truthfully, he wins the object and his payoff is 0. But if he bids 6, he increases his payoff to 1.

Bailey-Cavallo mechanism

Second-price auction is a natural approach in the set up when the central authority is a seller, as the tax corresponds then to payment for the object for sale. But we can also use the initial decision problem simply to determine which of the player values the object most. In such a set up the central authority is merely an arbiter and it is meaningful then to reach the decision with limited taxes.

Below, given a sequence $\theta \in \mathbb{R}^n$ of reals we denote by θ^* its reordering from the largest to the smallest element. So for example, for $\theta = (1, 4, 2, 3, 0)$ we have $(\theta_{-2})^*_2 = 2$ since $\theta_{-2} = (1, 2, 3, 0)$ and $(\theta_{-2})^* = (3, 2, 1, 0)$.

In the case of the second-price auction the final social welfare, i.e., $\sum_{j=1}^n u_j((f,t)(\theta), \theta_j)$, equals $\theta_i - \max_{j \neq i} \theta_j$, where $i = \text{argsmax}\theta$, so it equals the difference between the highest bid and the second highest bid.

We now discuss a modification of the second-price auction which yields a larger final social welfare. To ensure that it is well-defined we need to assume that $n \geq 3$. This modification, called ***Bailey-Cavallo mechanism***, is achieved by combining each tax $t'_i(\theta)$ to be paid in the second-price auction with

$$h'_i(\theta_{-i}) := \frac{(\theta_{-i})^*_2}{n},$$

that is, by using

$$t_i(\theta) := t'_i(\theta) + h'_i(\theta_{-i}).$$

Note that this yields a Groves mechanism since by the definition of the pivotal mechanism for specific functions $h_1, \ldots, h_n$

$$t'_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) + h_i(\theta_{-i}),$$

and consequently

$$t_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) + (h_i + h'_i)(\theta_{-i}).$$

In fact, this modification is a Groves mechanism if we start with an arbitrary Groves mechanism. In the case of the second-price auction the resulting mechanism is feasible since for all $i \in \{1, \ldots, n\}$ and θ we have $(\theta_{-i})^*_2 \leq \theta^*_2$ and as a result, since $\max_{j \neq i} \theta_j = \theta^*_2$,

$$\sum_{i=1}^{n} t_i(\theta) = \sum_{i=1}^{n} t'_i(\theta) + \sum_{i=1}^{n} h'_i(\theta_{-i}) = \sum_{i=1}^{n} \frac{-\theta^*_2 + (\theta_{-i})^*_2}{n} \leq 0.$$

Let, given the sequence θ of submitted bids (types), π be the permutation of $1, \ldots, n$ such that $\theta_{\pi(i)} = \theta^*_i$ for $i \in \{1, \ldots, n\}$ (where we break the ties by selecting players with the lower index first). So the ith highest bid is by player $\pi(i)$ and the object is sold to player $\pi(1)$. Note that then

- $(\theta_{-i})^*_2 = \theta^*_3$ for $i \in \{\pi(1), \pi(2)\}$,
- $(\theta_{-i})^*_2 = \theta^*_2$ for $i \in \{\pi(3), \ldots, \pi(n)\}$,

so the above mechanism boils down to the following payments by player $\pi(1)$:

- $\frac{\theta^*_3}{n}$ to player $\pi(2)$,
- $\frac{\theta^*_2}{n}$ to players $\pi(3), \ldots, \pi(n)$,
- $\theta^*_2 - \frac{2}{n}\theta^*_3 - \frac{n-2}{n}\theta^*_2 = \frac{2}{n}(\theta^*_2 - \theta^*_3)$ to the central authority.

To illustrate these payments assume that there are three players, A, B, and C whose true types (valuations) are 18, 21, and 24, respectively. When they bid truthfully the object is allocated to player C. In the second-price auction player's C tax is 21 and the final social welfare is $24 - 21 = 3$.

In constrast, in the case of the Bailey-Cavallo mechanism we have for the vector $\theta = (18, 21, 24)$ of submitted types $\theta^*_2 = 21$ and $\theta^*_3 = 18$, so player C pays

- 6 to player B,
- 7 to player A,
- 2 to the tax authority.

So the final social welfare is now $24 - 2 = 22$. Table 11.1 summarizes the situation.

player	type	tax	u_i
A	18	7	7
B	21	6	6
C	24	-15	9

Table 11.1: The Bailey-Cavallo mechanism

Public project problem

Let us return now to Example 22. To compute the taxes in the case of the pivotal mechanism we use the following observation.

Note 51 *In the public project problem we have for the pivotal mechanism*

$$t_i(\theta) = \begin{cases} 0 & \text{if } \sum_{j \neq i} \theta_j \geq \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j \geq c \\ \sum_{j \neq i} \theta_j - \frac{n-1}{n}c & \text{if } \sum_{j \neq i} \theta_j < \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j \geq c \\ 0 & \text{if } \sum_{j \neq i} \theta_j \leq \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j < c \\ \frac{n-1}{n}c - \sum_{j \neq i} \theta_j & \text{if } \sum_{j \neq i} \theta_j > \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j < c \end{cases}$$

□

To illustrate the pivotal mechanism suppose that there are three players, A, B, and C whose true types are 6, 7, and 25, and $c = 30$, respectively. When these types are announced the project takes place and Table 11.2 summarizes the taxes that players need to pay and their final utilities. The taxes were computed using Note 51.

player	type	tax	u_i
A	6	0	-4
B	7	0	-3
C	25	-7	8

Table 11.2: The pivotal mechanism for the public project problem

Suppose now that the true types of players are 4, 3 and 22, respectively and, as before, $c = 30$. When these types are also the announced types, the project does not take place. Still, some players need to pay a tax, as Table 11.3 illustrates.

player	type	tax	u_i
A	4	-5	-5
B	3	-6	-6
C	22	0	0

Table 11.3: The pivotal mechanism for the public project problem

Reversed sealed-bid auction

Note that the pivotal mechanism is not appropriate here. Indeed, we noted already that in the pivotal mechanism all players need to *make* a payment to the central authority, while in the context of the reversed sealed-bid auction we want to ensure that the lowest bidder *receives* a payment from the authority and other bidders neither pay nor receive any payment.

This can be realized by using the Groves mechanism with the following tax definition:

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D \setminus \{i\}} \sum_{j \neq i} v_j(d, \theta_j).$$

The crucial difference between this mechanism and the pivotal mechanism is that in the second expression we take a maximum over all decisions in the set $D \setminus \{i\}$ and *not* D.

To compute the taxes in the reversed sealed-bid auction with the above mechanism we use the following observation.

Note 52

$$t_i(\theta) = \begin{cases} -\max_{j \neq i} \theta_j & \text{if } i = \mathit{argsmax}\, \theta. \\ 0 & \text{otherwise} \end{cases}$$

□

This is identical to Note 50 in which the taxes for the pivotal mechanism for the sealed bid auction were computed. However, because we use here negative reals as bids the interpretation is different. Namely, the taxes are now positive, i.e., the players now *receive* the payments. More precisely, the winner, i.e., player i such that $i = \text{argsmax}\, \theta$, receives the payment equal to the second lowest offer, while the other players pay no taxes.

For example, when $\theta = (-8, -5, -4, -6)$, the service is bought from player 3 who submitted the lowest bid, namely 4. He receives for it the amount 5. Indeed, $3 = \text{argsmax}\, \theta$ and $-\max_{j \neq 3} \theta_j = -(-5) = 5$.

Buying a path in a network

As in the case of the reversed sealed-bid auction the pivotal mechanism is not appropriate here since we want to ensure that the players whose edge was selected *receive* a payment. Again, we achieve this by a simple modification of the pivotal mechanism. We modify it to a Groves mechanism in which

- the central authority is viewed as an agent who procures an $s - t$ path and pays the players whose edges are used,
- the players have an incentive to participate: if an edge is used, then the final utility of its owner is ≥ 0.

Recall that in the case of the pivotal mechanism we have

$$t'_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{p \in D(G)} \sum_{j \neq i} v_j(p, \theta_j),$$

where we now explicitly indicate the dependence of the decision set on the underlying graph, i.e., $D(G) := \{p \mid p \text{ is a } s - t \text{ path in } G\}$.

We now put instead

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{p \in D(G \setminus \{i\})} \sum_{j \neq i} v_j(p, \theta_j).$$

The following note provides the intuition for the above tax. We abbreviate here $\sum_{j \in p} \theta_j$ to $cost(p)$.

Note 53

$$t_i(\theta) = \begin{cases} cost(p_2) - cost(p_1 - \{i\}) & \text{if } i \in p_1 \\ 0 & \text{otherwise} \end{cases}$$

where p_1 is the shortest $s - t$ path in $G(\theta)$ and p_2 is the shortest $s - t$ path in $(G \setminus \{i\})(\theta_{-i})$.

Proof. Note that for each $s - t$ path p we have

$$-\sum_{j \neq i} v_j(p, \theta_j) = \sum_{j \in p - \{i\}} \theta_j.$$

Recall now that $f(\theta)$ is the shortest $s-t$ path in $G(\theta)$, i.e., $f(\theta) = p_1$. So $\sum_{j \neq i} v_j(f(\theta), \theta_j) = -cost(p_1 - \{i\})$.

To understand the second expression in the definition of $t_i(\theta)$ note that for each $p \in D(G \setminus \{i\})$, so for each $s-t$ path p in $G \setminus \{i\}$, we have

$$-\sum_{j \neq i} v_j(p, \theta_j) = \sum_{j \in p - \{i\}} \theta_j = \sum_{j \in p} \theta_j,$$

since the edge i does not belong to the path p. So $-\max_{p \in D(G \setminus \{i\})} \sum_{j \neq i} v_j(p, \theta_j)$ equals the length of the shortest $s-t$ path in $(G \setminus \{i\})(\theta_{-i})$, i.e., it equals $cost(p_2)$. $\square$

So given θ and the above definitions of the paths p_1 and p_2 the central authority pays to each player i whose edge is used the amount $cost(p_2) - cost(p_1 - \{i\})$. The final utility of such a player is then $-\theta_i + cost(p_2) - cost(p_1 - \{i\})$, i.e., $cost(p_2) - cost(p_1)$. So by the choice of p_1 and p_2 it is positive. No payments are made to the other players and their final utilities are 0.

Consider an example. Take the communication network depicted in Figure 11.1.

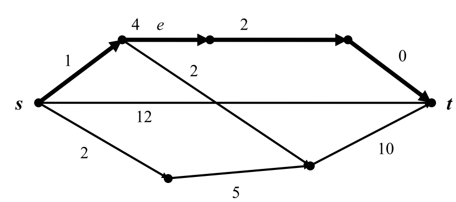

Figure 11.1: A communication network

This network has nine edges, so it corresponds to a decision problem with nine players. We assume that each player submitted the depicted length of the edge. Consider the player who owns the edge e, of length 4. To compute the payment he receives we need to determine the shortest $s-t$ path and the shortest $s-t$ path that does not include the edge e. The first path is the upper path, depicted in Figure 11.1 in bold. It contains the edge e and has the length 7. The second path is simply the edge connecting s and t and its length is 12. So, assuming that the players submit the costs truthfully, according to Note 53 player e receives the payment $12 - (7 - 4) = 9$ and his final utility is $9 - 4 = 5$.

11.4 Green and Laffont result

Until now we studied only one class of incentive compatible direct mechanisms, namely Groves mechanisms. Are there any other ones? J. Green and J.-J. Laffont showed that when the decision rule is efficient, under a natural assumption no other incentive compatible direct mechanisms exist. To formulate the relevant result we introduce the following notion.

Given a decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$, we call the utility function v_i ***complete*** if

$$\{v \mid v : D \to \mathbb{R}\} = \{v_i(\cdot, \theta_i) \mid \theta_i \in \Theta_i\},$$

that is, if each function $v : D \to \mathbb{R}$ is of the form $v_i(\cdot, \theta_i)$ for some $\theta_i \in \Theta_i$.

Theorem 54 *Consider a decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ with an efficient decision rule f. Suppose that each utility function v_i is complete. Then each incentive compatible direct mechanism is a Groves mechanism.*

To prove it first observe that each direct mechanism originating from a decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ can be written in a 'Groves-like' way, by putting

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) + h_i(\theta),$$

where each function h_i is defined on Θ and *not* on Θ_{-i}, as in the Groves mechanisms.

Lemma 55 *For each incentive compatible direct mechanism*

$$(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f, t)),$$

given the above representation, for all $i \in \{1, \ldots, n\}$

$$f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i}) \text{ implies } h_i(\theta_i, \theta_{-i}) = h_i(\theta'_i, \theta_{-i}).$$

Proof. Fix $i \in \{1, \ldots, n\}$. We have

$$u_i((f,t)(\theta_i, \theta_{-i}), \theta_i) = \sum_{j=1}^{n} v_j(f(\theta_i, \theta_{-i})), \theta_j) + h_i(\theta_i, \theta_{-i})$$

and

$$u_i((f,t)(\theta'_i, \theta_{-i}), \theta_i) = \sum_{j=1}^{n} v_j(f(\theta'_i, \theta_{-i})), \theta_j) + h_i(\theta'_i, \theta_{-i}),$$

so, on the account of the incentive compatibility, $f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i})$ implies $h_i(\theta_i, \theta_{-i}) \geq h_i(\theta'_i, \theta_{-i})$. By symmetry $h_i(\theta'_i, \theta_{-i}) \geq h_i(\theta_i, \theta_{-i})$, as well. □

Proof of Theorem 54.

Consider an incentive compatible direct mechanism

$$(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f,t))$$

and its 'Groves-like' representation with the functions $h_1, \ldots, h_n$. We need to show that no function h_i depends on its ith argument. Suppose otherwise. Then for some i, θ and θ'_i

$$h_i(\theta_i, \theta_{-i}) > h_i(\theta'_i, \theta_{-i}).$$

Choose an arbitrary ϵ from the open interval $(0, h_i(\theta_i, \theta_{-i}) - h_i(\theta'_i, \theta_{-i}))$ and consider the following function $v : D \rightarrow \mathbb{R}$:

$$v(d) := \begin{cases} \epsilon - \sum_{j \neq i} v_j(d, \theta_j) & \text{if } d = f(\theta'_i, \theta_{-i}) \\ -\sum_{j \neq i} v_j(d, \theta_j) & \text{otherwise} \end{cases}$$

By the completeness of v_i for some $\theta''_i \in \Theta_i$

$$v(d) = v_i(d, \theta''_i)$$

for all $d \in D$.

Since $h_i(\theta_i, \theta_{-i}) > h_i(\theta'_i, \theta_{-i})$, by Lemma 55 $f(\theta_i, \theta_{-i}) \neq f(\theta'_i, \theta_{-i})$, so by the definition of v

$$v_i(f(\theta_i, \theta_{-i}), \theta''_i) + \sum_{j \neq i} v_j(f(\theta_i, \theta_{-i}), \theta_j) = 0. \tag{11.1}$$

Further, for each $d \in D$ the sum $v_i(d, \theta_i'') + \sum_{j \neq i} v_j(d, \theta_j)$ equals either 0 or ϵ. This means that by the efficiency of f

$$v_i(f(\theta_i'', \theta_{-i}), \theta_i'') + \sum_{j \neq i} v_j(f(\theta_i'', \theta_{-i}), \theta_j) = \epsilon. \tag{11.2}$$

Hence, by the definition of v we have $f(\theta_i'', \theta_{-i}) = f(\theta_i', \theta_{-i})$, and consequently by Lemma 55

$$h_i(\theta_i'', \theta_{-i}) = h_i(\theta_i', \theta_{-i}). \tag{11.3}$$

We have now by (11.1)

$$\begin{aligned} & u_i((f,t)(\theta_i, \theta_{-i}), \theta_i'') \\ &= v_i(f(\theta_i, \theta_{-i}), \theta_i'') + \sum_{j \neq i} v_j(f(\theta_i, \theta_{-i}), \theta_j) + h_i(\theta_i, \theta_{-i}) \\ &= h_i(\theta_i, \theta_{-i}). \end{aligned}$$

In turn, by (11.2) and (11.3),

$$\begin{aligned} & u_i((f,t)(\theta_i'', \theta_{-i}), \theta_i'') \\ &= v_i(f(\theta_i'', \theta_{-i}), \theta_i'') + \sum_{j \neq i} v_j(f(\theta_i'', \theta_{-i}), \theta_j) + h_i(\theta_i'', \theta_{-i}) \\ &= \epsilon + h_i(\theta_i', \theta_{-i}). \end{aligned}$$

But by the choice of ϵ we have $h_i(\theta_i, \theta_{-i}) > \epsilon + h_i(\theta_i', \theta_{-i})$, so

$$u_i((f,t)(\theta_i, \theta_{-i}), \theta_i'') > u_i((f,t)(\theta_i'', \theta_{-i}), \theta_i''),$$

which contradicts the incentive compatibility for the joint type $(\theta_i'', \theta_{-i})$. $\square$

Chapter 12

Pre-Bayesian Games

Mechanism design, as introduced in the previous chapter, can be explained in game-theoretic terms using pre-Bayesian games In strategic games, after each player selected his strategy, each player knows the payoff of *every other player*. This is not the case in pre-Bayesian games in which each player has a private type on which he can condition his strategy. This distinguishing feature of pre-Bayesian games explains why they form a class of ***games with incomplete information***. Formally, they are defined as follows.

Assume a set $\{1, \ldots, n\}$ of players, where $n > 1$. A ***pre-Bayesian game*** for n players consists of

- a non-empty set A_i of ***actions***,
- a non-empty set Θ_i of ***types***,
- a ***payoff function*** $p_i : A_1 \times \ldots \times A_n \times \Theta_i \rightarrow \mathbb{R}$,

for each player i.

Let $A := A_1 \times \ldots \times A_n$. In a pre-Bayesian game Nature (an external agent) moves first and provides each player i with a type $\theta_i \in \Theta_i$. Each player knows only his type. Subsequently the players simultaneously select their actions. The payoff function of each player now depends on his type, so after all players selected their actions, each player knows his payoff but does not know the payoffs of the other players. Note that given a pre-Bayesian game, every joint type $\theta \in \Theta$ uniquely determines a strategic game, to which we refer below as a θ-game.

A ***strategy*** for player i in a pre-Bayesian game is a function $s_i : \Theta_i \rightarrow A_i$. A strategy $s_i(\cdot)$ for player i is called

- ***best response*** to the joint strategy $s_{-i}(\cdot)$ of the opponents of i if for all $a_i \in A_i$ and $\theta \in \Theta$

$$p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq p_i(a_i, s_{-i}(\theta_{-i}), \theta_i),$$

- ***dominant*** if for all $a \in A$ and $\theta_i \in \Theta_i$

$$p_i(s_i(\theta_i), a_{-i}, \theta_i) \geq p_i(a_i, a_{-i}, \theta_i),$$

Then a joint strategy $s(\cdot)$ is called an ***ex-post equilibrium*** if each $s_i(\cdot)$ is a best response to $s_{-i}(\cdot)$. Alternatively, $s(\cdot) := (s_1(\cdot), \ldots, s_n(\cdot))$ is an ex-post equilibrium if

$$\forall \theta \in \Theta \ \forall i \in \{1, \ldots, n\} \ \forall a_i \in A_i \ p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq p_i(a_i, s_{-i}(\theta_{-i}), \theta_i),$$

where $s_{-i}(\theta_{-i})$ is an abbreviation for the sequence of actions $(s_j(\theta_j))_{j \neq i}$.

So $s(\cdot)$ is an ex-post equilibrium iff for every joint type $\theta \in \Theta$ the sequence of actions $(s_1(\theta_1), \ldots, s_n(\theta_n))$ is a Nash-equilibrium in the corresponding θ-game. Further, $s_i(\cdot)$ is a dominant strategy of player i iff for every type $\theta_i \in \Theta_i$, $s_i(\theta_i)$ is a dominant strategy of player i in every (θ_i, θ_{-i})-game.

We also have the following immediate observation.

Note 56 (Dominant Strategy) *Consider a pre-Bayesian game G. Suppose that $s(\cdot)$ is a joint strategy such that each $s_i(\cdot)$ is a dominant strategy. Then it is an ex-post equilibrium of G.* □

Example 26 As an example of a pre-Bayesian game, suppose that

- $\Theta_1 = \{U, D\}$, $\Theta_2 = \{L, R\}$,
- $A_1 = A_2 = \{F, B\}$,

and consider the pre-Bayesian game uniquely determined by the following four θ-games. Here and below we marked the payoffs in Nash equilibria in these θ-games in bold.

L

U		F	B
	F	$\mathbf{2, 1}$	$2, 0$
	B	$0, 1$	$2, 1$

R

	F	B
F	$2, 0$	$\mathbf{2, 1}$
B	$0, 0$	$2, 1$

		F	B		F	B
D	F	$3,1$	$2,0$	F	$3,0$	$2,1$
	B	$\mathbf{5,1}$	$4,1$	B	$5,0$	$\mathbf{4,1}$

This shows that the strategies $s_1(\cdot)$ and $s_2(\cdot)$ such that

$$s_1(U) := F,\ s_1(D) := B,\ s_2(L) = F,\ s_2(R) = B$$

form here an ex-post equilibrium. □

However, there is a crucial difference between strategic games and pre-Bayesian games. We call a pre-Bayesian game ***finite*** if each set of actions and each set of types is finite. By the ***mixed extension*** of a finite pre-Bayesian game

$$(A_1, \ldots, A_n, \Theta_1, \ldots, \Theta_n, p_1, \ldots, p_n)$$

we mean below the pre-Bayesian game

$$(\Delta A_1, \ldots, \Delta A_n, \Theta_1, \ldots, \Theta_n, p_1, \ldots, p_n).$$

Example 27 Consider the following pre-Bayesian game:

- $\Theta_1 = \{U, B\}$, $\Theta_2 = \{L, R\}$,
- $A_1 = A_2 = \{C, D\}$,

		L			R	
		C	D		C	D
U	C	$2,2$	$0,0$	C	$2,1$	$0,0$
	D	$3,0$	$\mathbf{1,1}$	D	$3,0$	$\mathbf{1,2}$
		C	D		C	D
B	C	$\mathbf{1,2}$	$3,0$	C	$\mathbf{1,1}$	$3,0$
	D	$0,0$	$2,1$	D	$0,0$	$2,2$

Even though each θ-game has a Nash equilibrium, they are so ‘positioned’ that the pre-Bayesian game has no ex-post equilibrium. Even more, if we consider a mixed extension of this game, then the situation does not change. The reason is that no new Nash equilibria are then added to the original θ-games.

Indeed, each of these original θ-games is solved by IESDS and hence by the IESDMS Theorem 39(ii) has a unique Nash equilibrium. This shows that a mixed extension of a finite pre-Bayesian game does not need to have an ex-post equilibrium, which contrasts with the existence of Nash equilibria in mixed extensions of finite strategic games. □

This motivates the introduction of a new notion of an equilibrium. A strategy $s_i(\cdot)$ for player i is called ***safety-level best response*** to the joint strategy $s_{-i}(\cdot)$ of the opponents of i if for all strategies $s'_i(\cdot)$ of player i and all $\theta_i \in \Theta_i$

$$\min_{\theta_{-i}\in\Theta_{-i}} p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq \min_{\theta_{-i}\in\Theta_{-i}} p_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i).$$

Then a joint strategy $s(\cdot)$ is called a ***safety-level equilibrium*** if each $s_i(\cdot)$ is a safety-level best response to $s_{-i}(\cdot)$.

The following theorem was established by Monderer and Tennenholz.

Theorem 57 *Every mixed extension of a finite pre-Bayesian game has a safety-level equilibrium.* □

We now relate pre-Bayesian games to mechanism design. To this end we need one more notion. We say that a pre-Bayesian game is of a ***revelation-type*** if $A_i = \Theta_i$ for all $i \in \{1, \ldots, n\}$. So in a revelation-type pre-Bayesian game the strategies of a player are the functions on his set of types. A strategy for player i is called then ***truth-telling*** if it is the identity function $\pi_i(\cdot)$ on Θ_i.

Now mechanism design can be viewed as an instance of the revelation-type pre-Bayesian games. Indeed, we have the following immediate, yet revealing observation.

Theorem 58 *Given a direct mechanism*

$$(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f, t))$$

associate with it a revelation-type pre-Bayesian game, in which each payoff function p_i is defined by

$$p_i((\theta'_i, \theta_{-i}), \theta_i) := u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

Then the mechanism is incentive compatible iff in the associated pre-Bayesian game for each player truth-telling is a dominant strategy.

By Groves Theorem 49 we conclude that in the pre-Bayesian game associated with a Groves mechanism, $(\pi_1(\cdot), \ldots, \pi_n(\cdot))$ is a dominant strategy ex-post equilibrium.

Chapter 13

Alternative Concepts

In the presentation until now we heavily relied on the definition of a strategic game and focused several times on the crucial notion of a Nash equlibrium. However, both the concept of an equilibrium and of a strategic game can be defined in alternative ways. Here we discuss some alternative definitions and explain their consequences.

13.1 Other equilibria notions

Nash equilibrium is a most popular and most widely used notion of an equilibrium. However, there are many other natural alternatives. In this section we briefly discuss three alternative equilibria notions. To define them fix a strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.

Strict Nash equilibrium We call a joint strategy s a ***strict Nash equilibrium*** if

$$\forall i \in \{1, \ldots, n\} \; \forall s'_i \in S_i \setminus \{s_i\} \; p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i}).$$

So a joint strategy is a strict Nash equilibrium if each player achieves a *strictly lower* payoff by unilaterally switching to another strategy.

Obviously every strict Nash equilibrium is a Nash equilibrium and the converse does not need to hold.

Consider now the Battle of the Sexes game. Its pure Nash equilibria that we identified in Chapter 1 are clearly strict. However, its Nash equilibrium in mixed strategy we identified in Example 18 of Section 9.1 is not strict.

Indeed, the following simple observation holds.

Note 59 *Consider a mixed extension of a finite strategic game. Every strict Nash equilibrium is a Nash equilibrium in pure strategies.*

Proof. It is a direct consequence of the Characterization Lemma 28. □

Consequently each finite game with no Nash equilibrium in pure strategies, for instance the Matching Pennies game, has no strict Nash equilibrium in mixed strategies. So the analogue of Nash theorem does not hold for strict Nash equilibria, which makes this equilibrium notion less useful.

ϵ-Nash equilibrium The idea of an ϵ-Nash equilibrium formalizes the intuition that a joint strategy can be also be satisfactory for the players when each of them can gain only very little from deviating from his strategy.

Let $\epsilon > 0$ be a small positive real. We call a joint strategy s an ***ϵ-Nash equilibrium*** if

$$\forall i \in \{1, \ldots, n\} \; \forall s'_i \in S_i \; p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i}) - \epsilon.$$

So a joint strategy is an ϵ-Nash equilibrium if no player can gain more than ϵ by unilaterally switching to another strategy. In this context ϵ can be interpreted either as the amount of uncertainty about the payoffs or as the gain from switching to another strategy.

Clearly, a joint strategy is a Nash equilibrium iff it is an ϵ-Nash equilibrium for every $\epsilon > 0$. However, the payoffs in an ϵ-Nash equilibrium can be substantially lower than in a Nash equilibrium. Consider for example the following game:

	L	R
T	$1, \quad 1$	$0, \quad 0$
B	$1+\epsilon, \quad 1$	$100, 100$

This game has a unique Nash equilibrium (B, R), which obviously is also an ϵ-Nash equilibrium. However, (T, L) is also an ϵ-Nash equilibrium.

Strong Nash equilibrium Another variation of the notion of a Nash equilibrium focusses on the concept of a coalition, by which we mean a non-empty subset of all players.

Given a subset $K := \{k_1, \ldots, k_m\}$ of $N := \{1, \ldots, n\}$ we abbreviate the sequence $(s_{k_1}, \ldots, s_{k_m})$ of strategies to s_K and $S_{k_1} \times \ldots \times S_{k_m}$ to S_K.

We call a joint strategy s a ***strong Nash equilibrium*** if for all coalitions K there does not exist $s'_K \in S_K$ such that

$$p_i(s'_K, s_{N \setminus K}) > p_i(s_K, s_{N \setminus K}) \text{ for all } i \in K.$$

So a joint strategy is a strong Nash equilibrium if no coalition can profit from deviating from it, where by "profit from" we mean that each member of the coalition gets a strictly higher payoff. The notion of a strong Nash equilibrium generalizes the notion of a Nash equilibrium by considering possible deviations of coalitions instead of individual players.

Note that the unique Nash equilibrium of the Prisoner's Dilemma game is strict but not strong. For example, if both players deviate from D to C, then each of them gets a strictly higher payoff.

Correlated equilibrium The final concept of an equilibrium that we introduce is a generalization of Nash equilibrium in mixed strategies. Recall from Chapter 9 that given a finite strategic game $G := (S_1, \ldots, S_n, p_1, \ldots, p_n)$ each joint mixed strategy $m = (m_1, \ldots, m_n)$ induces a probability distribution over S, defined by

$$m(s) := m_1(s_1) \cdot \ldots \cdot m_n(s_n),$$

where $s \in S$.

We have then the following observation.

Note 60 (Nash Equilibrium in Mixed Strategies) *Consider a finite strategic game $(S_1, \ldots, S_n, p_1, \ldots, p_n)$.*

Then m is a Nash equilibrium m in mixed strategies iff for all $i \in \{1, \ldots, n\}$ and all $s'_i \in S_i$

$$\sum_{s \in S} m(s) \cdot p_i(s_i, s_{-i}) \geq \sum_{s \in S} m(s) \cdot p_i(s'_i, s_{-i}).$$

Proof. Fix $i \in \{1, \ldots, n\}$ and choose some $s'_i \in S_i$. Let

$$m'_i(s_i) := \begin{cases} 1 & \text{if } s_i = s'_i \\ 0 & \text{otherwise} \end{cases}$$

So m'_i is the mixed strategy that represents the pure strategy s'_i.

Let now $m' := (m_1, \ldots, m_{i-1}, m'_i, m_{i+1}, \ldots, m_n)$. We have

$$p_i(m) = \sum_{s \in S} m(s) \cdot p_i(s_i, s_{-i})$$

and

$$p_i(s'_i, m_{-i}) = \sum_{s \in S} m'(s) \cdot p_i(s_i, s_{-i}).$$

Further, one can check that

$$\sum_{s \in S} m'(s) \cdot p_i(s_i, s_{-i}) = \sum_{s \in S} m(s) \cdot p_i(s'_i, s_{-i}).$$

So the claim is a direct consequence of the equivalence between items (i) and (ii) of the Characterization Lemma 28. □

We now generalize the above inequality to an arbitrary probability distribution over S. This yields the following equilibrium notion. We call a probability distribution π over S a ***correlated equilibrium*** if for all $i \in \{1, \ldots, n\}$ and all $s'_i \in S_i$

$$\sum_{s \in S} \pi(s) \cdot p_i(s_i, s_{-i}) \geq \sum_{s \in S} \pi(s) \cdot p_i(s'_i, s_{-i}).$$

By the above Note every Nash equilibrium in mixed strategies is a correlated equilibrium. To see that the converse is not true consider the Battle of the Sexes game:

	F	B
F	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

It is easy to check that the following probability distribution forms a correlated equilibrium in this game:

	F	B
F	$\frac{1}{2}$	0
B	0	$\frac{1}{2}$

Intuitively, this equilibrium corresponds to a situation when an external observes flips a fair coin and gives each player a recommendation which strategy to choose.

Exercise 14 Check the above claim. □

13.2 Variations on the definition of strategic games

The notion of a strategic game is quantitative in the sense that it refers through payoffs to real numbers. A natural question to ask is: do the payoff values matter? The answer depends on which concepts we want to study. We mention here three qualitative variants of the definition of a strategic game in which the payoffs are replaced by preferences. By a ***preference relation*** on a set A we mean here a linear ordering on A.

In Osborne and Rubinstein [1994] a strategic game is defined as a sequence

$$(S_1, \ldots, S_n, \succeq_1, \ldots, \succeq_n),$$

where each $\succeq_i$ is player's i *preference relation* defined on the set $S_1 \times \ldots \times S_n$ of joint strategies.

In Apt, Rossi and Venable [2008] another modification of strategic games is considered, called a ***strategic game with parametrized preferences***. In this approach each player i has a non-empty set of strategies S_i and a *preference relation* $\succeq_{s_{-i}}$ on S_i *parametrized* by a joint strategy s_{-i} of his opponents. In Apt, Rossi and Venable [2008] only strict preferences were considered and so defined finite games with parametrized preferences were compared with the concept of ***CP-nets*** (Conditional Preference nets), a formalism used for representing conditional and qualitative preferences, see, e.g., Boutilier et al. [2004].

Next, in Roux, Lescanne and Vestergaard [2008] ***conversion/preference games*** are introduced. Such a game for n players consists of a set S of *situations* and for each player i a *preference relation* $\succeq_i$ on S and a *conversion relation* $\rightarrow_i$ on S. The definition is very general and no conditions are placed on the preference and conversion relations. These games are used to formalize gene regulation networks and some aspects of security.

Another generalization of strategic games, called ***graphical games***, introduced in Kearns, Littman and Singh [2001]. These games stress the locality in taking decision. In a graphical game the payoff of each player depends only on the strategies of its neighbours in a given in advance graph structure over the set of players. Formally, such a game for n players with the corresponding strategy sets $S_1, \ldots, S_n$ is defined by assuming a neighbour function N that given a player i yields its set of neighbours $N(i)$. The payoff for player i is then a function p_i from $\times_{j \in N(i) \cup \{i\}} S_j$ to $\mathbb{R}$.

In all mentioned variants it is straightforward to define the notion of a Nash equilibrium. For example, in the conversion/preferences games it is defined as a situation s such that for all players i, if $s \rightarrow_i s'$, then $s' \not\succ_i s$. However, other introduced notions can be defined only for some variants. In particular, Pareto efficiency cannot be defined for strategic games with parametrized preferences since it requires a comparison of two arbitrary joint strategies. In turn, the notions of dominance cannot be defined for the conversion/preferences games, since they require the concept of a strategy for a player.

Various results concerning finite strategic games, for instance the IESDS Theorem 2, carry over directly to the the strategic games as defined in Osborne and Rubinstein [1994] or in Apt, Rossi and Venable [2008]. On the other hand, in the variants of strategic games that rely on the notion of a preference we cannot consider mixed strategies, since the outcomes of playing different strategies by a player cannot be aggregated.

References

K. R. Apt, F. Rossi, and K. B. Venable

[2008] Comparing the notions of optimality in CP-nets, strategic games and soft constraints, *Annals of Mathematics and Artificial Intelligence*, 52, pp. 25–54.

K. Binmore

[1991] *Fun and Games: A Text on Game Theory*, D.C. Heath.

[2007] *Playing for Real: A Text on Game Theory*, Oxford University Press, Oxford.

C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole

[2004] CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements., *J. Artif. Intell. Res. (JAIR)*, 21, pp. 135–191.

D. Fudenberg and J. Tirole

[1991] *Game Theory*, MIT Press, Cambridge, Massachusetts.

G. Jehle and P. Reny

[2000] *Advanced Microeconomic Theory*, Addison Wesley, Reading, Massachusetts, second ed.

M. Kearns, M. Littman, and S. Singh

[2001] Graphical models for game theory, in: *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI '01)*, Morgan Kaufmann, pp. 253–260.

A. Mas-Collel, M. D. Whinston, and J. R. Green

[1995] *Microeconomic Theory*, Oxford University Press, Oxford.

R. B. Myerson

[1991] *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Massachusetts.

J. F. Nash

[1950] Equilibrium points in n-person games, *Proceedings of the National Academy of Sciences, USA*, 36, pp. 48–49.

[1951] Non-cooperative games, *Annals of Mathematics*, 54, pp. 286–295.

N. Nisan, T. Roughgarden, É. Tardos, and V. J. Vazirani

[2007] eds., *Algorithmic Game Theory*, Cambridge University Press.

M. J. Osborne

[2005] *An Introduction to Game Theory*, Oxford University Press, Oxford.

M. J. Osborne and A. Rubinstein

[1994] *A Course in Game Theory*, The MIT Press, Cambridge, Massachusetts.

H. Peters

[2008] *Game Theory: A Multi-Leveled Approach*, Springer, Berlin.

K. Ritzberger

[2002] *Foundations of Non-cooperative Game Theory*, Oxford University Press, Oxford.

S. L. Roux, P. Lescanne, and R. Vestergaard

[2008] Conversion/preference games, *CoRR*, abs/0811.0071.

Y. Shoham and K. Leyton-Brown

[2008] *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*, Morgan and Claypool Publishers, Princeton.

S. Tijs

[2003] *Introduction to Game Theory*, Hindustan Book Agency, Gurgaon, India.

J. von Neumann and O. Morgenstern

[1944] *Theory of Games and Economic Behavior*, Princeton University Press.